arXiv:1510.08121v1 [cs.PL] 27 Oct 2015

TYPE DIRECTED SYNTHESIS OF PRODUCTS

Jonathan Frankle

A THESIS

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF MASTER OF SCIENCE IN ENGINEERING

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF

COMPUTER SCIENCE

Adviser: David Walker

September 2015

© Copyright by Jonathan Frankle, 2015. All rights reserved.

Abstract

Software synthesis - the process of generating complete, general-purpose programs from specifications - has become a hot research topic in the past few years. For decades the problem was thought to be insurmountable: the search space of possible programs is far too massive to efficiently traverse. Advances in efficient constraint solving have overcome this barrier, enabling a new generation of effective synthesis systems. Most existing systems compile synthesis tasks down to low-level SMT instances, sacrificing high-level semantic information while solving only first-order problems (i.e., filling integer holes). Recent work takes an alternative approach, using the Curry-Howard isomorphism and techniques from automated theorem proving to construct higher-order programs with algebraic datatypes.

My thesis involved extending this *type-directed* synthesis engine to handle product types, which required significant modifications to both the underlying theory and the tool itself. Product types streamline other language features, eliminating variable-arity constructors among other workarounds employed in the original synthesis system. A form of logical conjunction, products are *invertible*, making it possible to equip the synthesis system with an efficient theorem-proving technique called *focusing* that eliminates many of the nondeterministic choices inherent in proof search. These theoretical enhancements informed a new version of the type-directed synthesis prototype implementation, which remained performance-competitive with the original synthesizer. A significant advantage of the type-directed synthesis framework is its extensibility; this thesis is a roadmap for future such efforts to increase the expressive power of the system.

1 Introduction

Since the advent of computer programming, the cycle of writing, testing, and debugging code has remained a tedious and error-prone undertaking. There is no middle ground between a program that is correct and one that is not, so software development demands often-frustrating levels of precision. The task of writing code is mechanical and repetitive, with boilerplate and common idioms consuming ever more valuable developer-time. Aside from the languages and platforms, the process of software engineering has changed little over the past four decades.

The discipline of software synthesis responds to these deficiencies with a simple question: if computers can automate so much of our everyday lives, why can they not do the same for the task of developing software? Outside of a few specialized domains, this question – until recently – had an equally simple answer: the search space of candidate programs is too large to explore efficiently. As we increase the number of abstract syntax tree (AST) nodes that a program might require, the search space explodes combinatorially. Synthesizing even small programs seems hopeless.

In the past decade, however, several research program analysis and synthesis systems have overcome these barriers and developed into useful programming tools. At the core of Sketch [13], Rosette [22, 23], and Leon [9] are efficient SAT and SMT-solvers. These tools automate many development tasks – including test case generation, verification, angelic nondeterminism, and synthesis [22] – by compiling programs into constraint-solving problems.

These synthesis techniques, while a major technological step forward, are still quite limited. Constraint-solving tasks are fundamentally first-order: although they efficiently fill integer and boolean holes, they cannot scale to higher-order programs. Furthermore, in the process of compiling synthesis problems into the low-level language of constraint-solvers, these methods sacrifice

high-level semantic and type information that might guide the synthesis procedure through the search space more efficiently.

Based on these observations, recent work by Steve Zdancewic and Peter-Michael Osera at the Univerity of Pennsylvania [14] explores a different approach to synthesis: theorem-proving. The Curry-Howard isomorphism permits us to treat the type of a desired program as a theorem whose proof is the synthesized program. Translating this idea into an algorithm, we can search for the solution to the synthesis problem using existing automated theorem-proving techniques. Since many programs inhabit the same type, users also provide input-output examples to better specify the desired function and constrain the synthesizer's result.

This *type-directed* approach scales to higher-order programs and preserves the high-level program structures that guide the synthesis process. The synthesis algorithm can be designed to search only for well-typed programs in normal form, drastically reducing the search space of possible ASTs. At the time of writing, both of these features are unique to type-directed synthesis.

Not only is this technique efficient, but it also has the ability to scale to language features that have, until now, remained beyond the reach of synthesis. Many desirable features, like file input-output, map to existing systems of logic that, when integrated into the synthesizer, instantly enable it to generate the corresponding programs. For example, one might imagine synthesizing effectful computation using monads by performing proof search in lax logic [16]. When these logical constructs are kept orthogonal to one another, features can be added and removed from the synthesis language in modular fashion. This extensibility is one of the most noteworthy benefits of type-directed synthesis.

This thesis represents the first such extension to the type-directed synthesis system. My research involved adding product types to the original synthesis framework, which until then captured only the simply typed lambda calculus with recursion and algebraic datatypes. In the process, I heavily revised and expanded both the formal judgments that govern the synthesis procedure and the code that implements it, integrating additional theorem proving techniques that pave the way for future language extensions.

Contributions of this thesis.

1. An analysis of the effect of searching only for typed programs or typed programs in normal form on the size of the number of programs at a particular type.
2. A revised presentation of the judgments for the original type-directed synthesis system.
3. An extension of the original synthesis judgments to include product types and the *focusing* technique.
4. Theorems about the properties of the updated synthesis judgments and focusing, including proofs of admissibility of focusing and soundness of the full system.
5. Updates to the type-directed synthesis prototype that implement the new judgments.
6. An evaluation of the performance of the updated implementation on several canonical programs involving product types.
7. A thorough survey of related work.
8. A discussion of future research avenues for type-directed synthesis with emphasis on the observation that examples are refinement types.

Overview. The body of this thesis is structured as follows: I begin with an in-depth look at the original type-directed synthesis system in Section 2. In Section 3, I introduce the theory underlying

synthesis of products and extensively discuss the focusing technique, which handles product types efficiently. I describe implementation changes made to add tuples to the synthesis framework and evaluate the performance of the modified system in Section 4. In Section 5, I discuss related work, including extended summaries of other synthesis systems. Finally, I outline future research directions in Section 6, with particular emphasis on using intersection and refinement types that integrate input-output examples into the type system. I conclude in Section 7. Proofs of theorems presented throughout this thesis appear in Appendix A.

2 Type-Directed Synthesis

2.1 Overview

The following is a brief summary of type-directed synthesis [14], presented loosely within the *Syntax-Guided Synthesis* [3] framework.

Background theory. The system synthesizes over the simply-typed lambda calculus with recursive functions and non-polymorphic algebraic datatypes. In practice, it uses a subset of OCaml with the aforementioned features. In order to guarantee that all programs terminate, the language permits only structural recursion. Pre-defined functions can be made available to the synthesis process if desired (i.e., providing *map* and *fold* to a synthesis task involving list manipulation).

Synthesis problem. A user specifies the name and type signature of a function to be synthesized. No additional structural guidance or "sketching" is provided.

Solution specification. The function's type information combined with input-output examples constrain the function to be synthesized. We can treat this set of examples as a *partial function* that we wish to generalize into a total function. Since type-directed synthesis aims to scale to multi-argument, higher-order functions, these examples can map multiple inputs, including other functions, to a single output. Functions are not permissible as output examples, however, since the synthesis procedure must be able to decidably test outputs for equality.

Optimality criterion. The synthesis problem as currently described lends itself to a very simple algorithm: create a function that (1) on an input specified in an example, supplies the corresponding output and (2) on all other inputs, returns an arbitrary, well-typed value. To avoid generating such useless, *over-fitted* programs, type-directed synthesis requires some notion of a *best* result. In practice, we create the smallest (measured in AST nodes) program that satisfies the specification, since a more generic, recursive solution will lead to a smaller program than one that merely matches on examples.

Search strategy. Type-directed synthesis treats the signature of the function in the synthesis problem as a theorem to be proved and uses a modified form of reverse proof search [15] that integrates input-output examples to generate a program. By the Curry-Howard isomorphism, a proof of the theorem is a program with the desired type; if it satisfies the input-output examples, then it is a solution to the overall synthesis problem.

Style. Type-directed synthesis can be characterized as a hybrid algorithm that has qualities of both deductive and inductive synthesis. It uses a set of rules to extract the structure of the input-output examples into a program, which is reminiscent of deductive synthesis algorithms that use a similar process on complete specifications. When guessing variables and function applications, however, type-directed synthesis generates terms and checks whether they satisfy the examples, an approach in the style of inductive strategies like CEGIS [20].

2.2 Case Study: List *length*

Before delving into the technical details of the theory, consider the process of synthesizing the list *length* function as illustrated in Figure 1.

In Figure 1a, we begin with our synthesis problem: a function with its type signature and a list of input-output examples. The `?` preceding each example represents the hole to which the example corresponds. Initially, the entire function to be synthesized comprises a single hole.

Enclosed within angle brackets is each example *world*, which we define as a pair of (1) bindings of names to values and (2) the *goal* value that should be taken on by the expression chosen to fill the hole when the names are bound to those values. For example, we can read the expression

$$?_1\texttt{:}\quad \langle \texttt{x = 1, y = 2; 5} \rangle$$

as

When x = 1 and y = 2, the expression synthesized to fill hole $?_1$ should evaluate to 5.

In our initial synthesis problem in Figure 1a, no names have yet been bound to values. We could imagine providing the synthesis instance with a library of existing functions, like *fold* and *map*, in which case our example worlds would contain already-bound names at the start of the synthesis process.

Each goal value is expressed as a partial function mapping input arguments to an output. For example,

$$\texttt{[1; 2] => inc => [2; 3]}$$

means that, on inputs `[1; 2]` and the increment function, the desired output is `[2; 3]`.

In Figure 1b, we observe that every example is a partial function mapping a `natlist` to a `nat` and, as such, we synthesize a function with a `natlist` argument called `ls`. We must update our examples in kind. Where before we had an example world of the form

$$\langle \cdot\texttt{; [4; 3] => 2} \rangle$$

we now extract the value of `ls` and move it to the list of names bound to values:

$$\langle \texttt{ls = [4; 3]; 2} \rangle$$

Since we have now synthesized a function, our hole is of type `nat` and the goal value is the remainder of the partial function with `ls` removed.

In Figure 1c, we synthesize a `match` statement to break down the structure of `ls`. This creates two holes, one for each branch of the `match` statement. We partition our set of examples between the two holes depending on the constructor of the value bound to `ls`. Those examples for which `ls = []` constrain the hole for the `Nil` branch; all other examples constrain the `Cons` branch.

In Figure 1d, we turn our attention to the `Nil` branch. Since we only have a single example, it is safe to simply synthesize the value in the example's goal position (namely the constructor `O`). In the other branch, too, every example has the same constructor (`S`), which we then generate in our

Description	(a) Initial synthesis problem.	(b) Synthesize a function.
Program	`len : natlist -> nat = ?`	`len (ls : natlist) : nat = ?`

Examples	`?:⟨·; [] => 0⟩`
	`?:⟨·; [3] => 1⟩`
	`?:⟨·; [4; 3] => 2⟩`
		`?:⟨len = ..., ls = [];  0⟩`
		`?:⟨len = ..., ls = [3]; 1⟩`
		`?:⟨len = ..., ls = [4; 3]; 2⟩`

Judgment		IREFINE-FIX

(c) Synthesize a match statement.	(d) Complete the `Nil` branch with the `O` constructor.

`len (ls : natlist) : nat =`
`  match ls with`
`  | Nil -> ?1`
`  | Cons(hd, tl) -> ?2`
	`len (ls : natlist) : nat =`
	`  match ls with`
	`  | Nil -> O`
	`  | Cons(hd, tl) -> ?2`

`?1:⟨len = ..., ls = [] ; 0⟩`
`?2:⟨len = ..., ls = [3], hd = 3, tl = []; 1⟩`
`?2:⟨len = ..., ls = [4; 3], hd = 4, tl = [3]; 2⟩`
	`?2:⟨len = ..., ls = [3], hd = 3, tl = []; 1⟩`
	`?2:⟨len = ..., ls = [4; 3], hd = 4, tl = [3]; 2⟩`

IREFINE-MATCH, EGUESS-CTX	IREFINE-CTOR

(e) Synthesize constructor `S` in the remaining branch.	(f) Synthesize an application to fill the final hole.

`len (ls : natlist) : nat =`
`  match ls with`
`  | Nil -> O`
`  | Cons(hd, tl) -> S(?2)`
	`len (ls : natlist) : nat =`
	`  match ls with`
	`  | Nil -> O`
	`  | Cons(hd, tl) -> S(len tl)`

`?2:⟨len = ..., ls = [3], hd = 3, tl = []; 0⟩` `?2:⟨len = ..., ls = [4; 3], hd = 4, tl = [3]; 1⟩`	
IREFINE-CTOR	IREFINE-GUESS, EGUESS-APP, EGUESS-CTX

Figure 1: A step-by-step derivation of the list length function in type-directed synthesis. A ? character refers to a hole in the program that the synthesis algorithm aims to fill. Each example *world*, delimited with ⟨ and ⟩, contains variable bindings to the left of the ; and the goal value to the right. The preceding ? indicates the hole to which the example world corresponds. For brevity, we write unary numbers in their Arabic equivalents (S (S (O)) is abbreviated as 2) and lists in equivalent OCaml syntax (Cons(2, Cons(1, Nil)) is [2; 1]). The names of all recursive functions in scope (in this case `len`) are always available in the list of variable bindings. They are bound to the partial functions comprising their definitions. The example for `len` is the initial partial function example at the beginning of the synthesis process; it is elided from the example worlds for space.

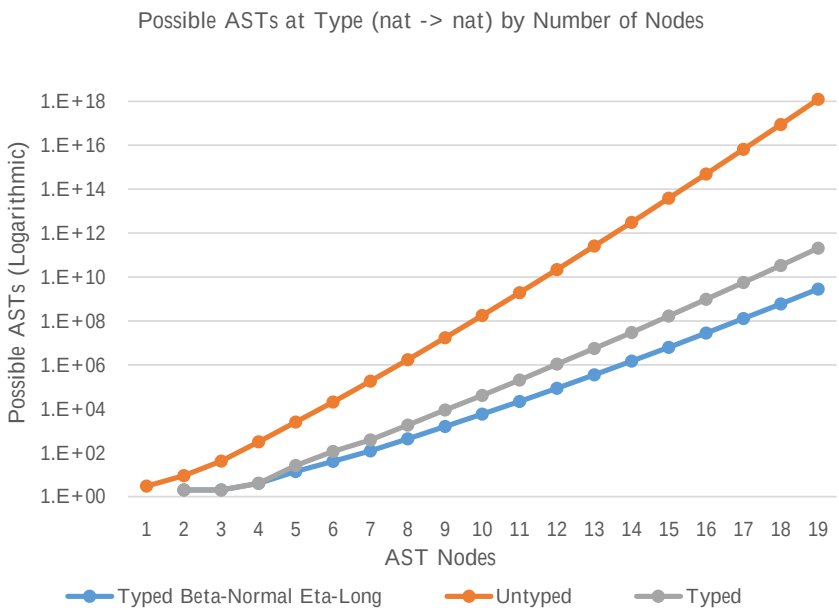

Figure 2: The number of possible ASTs (at type nat $\rightarrow$ nat) with a particular number of nodes. ASTs were generated over the type-directed synthesis language (the lambda calculus with recursive functions and algebraic datatypes). The top line (orange) includes all possible ASTs, while the middle and bottom lines include only typed (grey) and typed beta-normal eta-long (blue) ASTs respectively. The scale on the vertical axis is logarithmic.

program (Figure 1e). We update the value in the goal position accordingly, removing one use of the `S` constructor. Finally, with a recursive call to `len` on `tl`, our function is complete.

It is important to note that every function we generate is recursive. Therefore, the name of each function, including the top-level function of the synthesis problem, is available in every example world for which it is in scope. Function names are bound to the partial functions comprising their examples. Making these names available allows for recursive function calls. The examples for `len` are elided from example worlds in Figure 1 for readability.

2.3 Proof Search

Search strategy. The primary distinguishing quality of type-directed synthesis is its search strategy. The grammar of expressions e in the background theory is specified in Figure 3. Within this grammar are myriad ASTs that are not well-typed. As Figure 2 illustrates, merely restricting our search to well-typed terms drastically decreases the size of the search space. Even amongst well-typed terms, there are numerous expressions that are functionally identical. For example, $(\text{fix } f(x:\tau_1):\tau_1 = x)\ 1$ beta-reduces to 1 and $\text{fix } f(x:\tau_1):\tau_2 = (g\ x)$ eta-reduces to g. If we confine our search to programs in eta-long, beta-normal form, we avoid considering some of these duplicates and further reduce our search space. Figure 2 demonstrates the benefit of this further restriction: an additional one to two orders of magnitude reduction in the size of the search space.

Based on these observations, type-directed syntax avoids searching amongst all expressions (e), instead generating well-typed ASTs over a more restrictive grammar: that of introduction (I) and elimination (E) forms. Doing so guarantees that programs synthesized are in beta-normal form,

$$
\begin{array}{lll}
\tau & ::= \mathcal{B} \mid \tau_1 \to \tau_2 & \text{(Types: base types and function types)} \\
v & ::= C\ (v_1, ..., v_m) & \text{(Values: constructors, functions, and partial functions)} \\
 & \quad \mid \mathsf{fix}\ f(x:\tau_1):\tau_2 = e \mid pf & \\
pf & ::= \overline{v_i \Rightarrow \varepsilon_i}^{\,i\in m} & \text{(Partial functions: a finite map from inputs to outputs)} \\
\varepsilon & ::= C\ (\varepsilon_1, ..., \varepsilon_m) \mid pf & \text{(Examples)} \\
e & ::= x \mid C\ (e_1, ..., e_m) \mid pf & \text{(Expressions)} \\
 & \quad \mid \mathsf{fix}\ f(x:\tau_1):\tau_2 = e \mid e_1\ e_2 & \\
 & \quad \mid \mathsf{match}\ e\ \mathsf{with}\ \overline{C_i\ (x_1, ..., x_m) \to e_i}^{\,i\in l} & \\
 & & \\
E & ::= x \mid E\ I & \text{(Elimination forms)} \\
I & ::= E \mid C\ (I_1, ..., I_m) & \text{(Introduction forms)} \\
 & \quad \mid \mathsf{fix}\ f(x:\tau_1):\tau_2 = I & \\
 & \quad \mid \mathsf{match}\ E\ \mathsf{with}\ \overline{C_i\ (x_1, ..., x_m) \to I_i}^{\,i\in l} & \\
 & & \\
\Sigma & ::= \cdot \mid \Sigma, C : \tau_1 * ... * \tau_n \to \mathcal{B} & \text{(Constructor context: a map from constructor names to types)} \\
\Gamma & ::= \cdot \mid \Gamma, x:\tau & \text{(Type context: bindings of names to types)} \\
\mathcal{E} & ::= \cdot \mid \mathcal{E}, x \triangleright \varepsilon & \text{(Example context: names refined by examples)} \\
w & ::= \langle \mathcal{E}; \varepsilon \rangle & \text{(A single world: an example context and result example)} \\
\mathcal{W} & ::= \cdot \mid \mathcal{W}, w & \text{(A set of worlds)}
\end{array}
$$

Figure 3: Grammars for the synthesis judgments of type-directed synthesis.

$$\boxed{\Sigma \mid \Gamma \mid \mathcal{W} \vdash \tau \overset{I}{\rightsquigarrow} I}$$

IREFINE-GUESS

$$\frac{\Sigma \mid \Gamma \vdash \mathcal{B} \overset{E}{\rightsquigarrow} E \qquad \forall i \in n,\ \mathcal{E}_i(E) \to^* \varepsilon_i}{\Sigma \mid \Gamma \mid \overline{\langle \mathcal{E}_i; \varepsilon_i \rangle}^{\,i\in n} \vdash \mathcal{B} \overset{I}{\rightsquigarrow} E}$$

IREFINE-CTOR

$$\frac{C : \tau_1 * ... * \tau_m \to \mathcal{B} \in \Sigma \qquad \forall j \in m, \Sigma \mid \Gamma \mid \overline{\langle \mathcal{E}_i; \varepsilon_{(i,j)} \rangle}^{\,i\in n} \vdash \tau_j \overset{I}{\rightsquigarrow} I_j}{\Sigma \mid \Gamma \mid \overline{\langle \mathcal{E}_i; C\ (\varepsilon_{(i,1)}, ..., I_{(i,m)}) \rangle}^{\,i\in n} \vdash \mathcal{B} \overset{I}{\rightsquigarrow} C\ (I_1, ..., I_m)}$$

IREFINE-FIX

$$\frac{\Sigma \mid \Gamma, f : \tau_1 \to \tau_2, x : \tau_1 \mid \overline{\langle \mathcal{E}_i, f \triangleright \overline{v_{(i,j)} \Rightarrow \varepsilon_{(i,j)}}^{\,j\in m_i}, x \triangleright v_{(i,k)}; \varepsilon_{(i,k)} \rangle}^{\,(i,k)\in(n,m_i)} \vdash \tau_2 \overset{I}{\rightsquigarrow} I}{\Sigma \mid \Gamma \mid \overline{\langle \mathcal{E}_i; \overline{v_{(i,j)} \Rightarrow \varepsilon_{(i,j)}}^{\,j\in m_i} \rangle}^{\,i\in n} \vdash \tau_1 \to \tau_2 \overset{I}{\rightsquigarrow} \mathsf{fix}\ f(x:\tau_1):\tau_2 = I}$$

IREFINE-MATCH

$$\frac{\Sigma \mid \Gamma \vdash \mathcal{B} \overset{E}{\rightsquigarrow} E \qquad \forall j \in m,\ C_j : \tau_{(j,1)} * ... * \tau_{(j,l)} \to \mathcal{B} \in \Sigma \qquad \forall j \in m,\ \Sigma \mid \Gamma, x : \tau_{(j,1)}, ..., x : \tau_{(j,l)} \mid \{\langle \mathcal{E}_i, x_1 \triangleright v_1, ..., x_l \triangleright v_l; \varepsilon_i \rangle \mid \mathcal{E}_i(E) \to^* C_j\ (v_1, ..., v_l)\} \vdash \tau \overset{I}{\rightsquigarrow} I_j}{\Sigma \mid \Gamma \mid \overline{\langle \mathcal{E}_i; \varepsilon_i \rangle}^{\,i\in n} \vdash \tau \overset{I}{\rightsquigarrow} \mathsf{match}\ E\ \mathsf{with}\ \overline{C_j\ (x_1, ..., x_l) \to I_j}^{\,j\in m}}$$

$$\boxed{\Sigma \mid \Gamma \vdash \tau \overset{E}{\rightsquigarrow} E}$$

EGUESS-CTX

$$\frac{x : \tau \in \Gamma}{\Sigma \mid \Gamma \vdash \tau \overset{E}{\rightsquigarrow} x}$$

EGUESS-APP

$$\frac{\Sigma \mid \Gamma \vdash \tau_1 \to \tau \overset{E}{\rightsquigarrow} E \qquad \Sigma \mid \Gamma \mid \cdot \vdash \tau_1 \overset{I}{\rightsquigarrow} I}{\Sigma \mid \Gamma \vdash \tau \overset{E}{\rightsquigarrow} E\ I}$$

Figure 4: Synthesis judgments for type-directed synthesis. An expression with a bar and an index above it represents a set of expressions sharing the same structure.

since no lambda term can appear on the left side of an application. If, in addition, we generate only lambdas at function type, programs will also be eta-long.

We can combine these ideas into a proof search algorithm that begins with a synthesis problem as in Figure 1 and uses type information, along with the structure of examples, to generate programs. The synthesis judgments for this algorithm appear in Figure 4.

Refining introduction forms. The judgment for producing introduction forms, which is also the top-level synthesis judgment, is of the form

$$\Sigma \mid \Gamma \mid \mathcal{W} \vdash \tau \overset{I}{\rightsquigarrow} I$$

which states:

Given constructors Σ and names bound to types Γ, we synthesize introduction form I at type τ conforming to the examples in worlds $\mathcal{W}$.

The IREFINE-FIX rule extracts the structure of a partial function example as in Figure 1b, synthesizing a fixpoint and creating a new subproblem for the fixpoint's body that can refer to the argument and make recursive calls. The IREFINE-CTOR rule observes that every example shares the same constructor and therefore synthesizes the constructor in the program, likewise creating subproblems for each of the constructor's arguments.

The IREFINE-MATCH rule generates a `match` statement, guessing an elimination-form scrutinee at base type on which to pattern-match. The rule creates a sub-problem for each branch of the match statement - that is, one for every constructor of the scrutinee's type. The existing set of examples is partitioned among these sub-problems according to the value that the scrutinee takes on when evaluated in each example world. Since the values of the bound variables may vary from world to world, the scrutinee can evaluate to different constructors in different contexts. This evaluation step determines which examples correspond to which branches, allowing us to constrain the sub-problems and continue with the synthesis process.

Finally, at base type ($\mathcal{B}$), we may also guess an elimination form, captured by the IREFINE-GUESS rule. The restriction that this only occurs at base type helps enforce the property that we only generate eta-long programs.

Guessing elimination forms. The judgment form for guessing elimination forms

$$\Sigma \mid \Gamma \vdash \tau \overset{E}{\rightsquigarrow} E$$

states:

Given constructors Σ and names bound to types Γ, we synthesize elimination form E at type τ.

Observe that this judgment form does not include examples. Unlike the introduction forms, we cannot use the structural content of the examples to inform the elimination forms we guess. Instead, we may only ensure that the elimination form that we generate obeys the examples. This requirement appears in the additional condition of the IREFINE-GUESS rule

$$\forall i \in n,\ \mathcal{E}_i(E) \rightarrow^* \varepsilon_i$$

reading:

For each of our n examples, the elimination form E must, when substituted with the value bindings in $\mathcal{E}_i$, evaluate to the corresponding goal value ε_i.

The EGUESS-APP rule guesses a function application by recursively guessing a function and an argument. The function must be an elimination form, which ensures that we only guess beta-normal programs. The argument may be an introduction form, but the call back into an IREFINE rule does so without examples. Finally, the base case EGUESS-CTX guesses a name from the context Γ.

Nondeterminism. These rules are highly nondeterministic: at any point, several rules apply. This is especially pertinent at base type, where we could refine a constructor or match statement while guessing an elimination form. Figure 1 traces only one of many paths through the expansive search space that this nondeterminism creates.

Nondeterminism provides the flexibility for an optimality condition (Section 2.1) to select the best possible program. Although we could continually apply IREFINE rules until we fully extract the structure of the examples, doing so would generate a larger than necessary program that has been *overfitted* to the examples. Instead, the EGUESS rules allow us to make recursive calls that create smaller, more general programs.

Relationship with other systems. The synthesis judgments are closely related to two other systems: the sequent calculus and bidirectional type-checking. We can reframe the judgment

$$\Sigma \mid x_1 : \tau_1, ..., x_n : \tau_n \mid \mathcal{W} \vdash \tau \overset{I}{\rightsquigarrow} I$$

as the sequent

$$\tau_1, ..., \tau_n \Longrightarrow \tau$$

if we elide the proof terms (x_1 through x_n and I) and examples. Many of our rules (particularly IREFINE-FIX and EGUESS-APP) directly mirror their counterparts in Gentzen's characterization of the sequent calculus [6]. The only modifications are the addition of proof terms and bookkeeping to manage examples. The rules that handle algebraic datatypes and pattern matching (IREFINE-CTOR and IREFINE-MATCH) loosely resemble logical disjunction, which Gentzen also describes.

The type-directed synthesis system also resembles bidirectional typechecking. Where traditional bidirectional typechecking rules are functions that synthesize or check a type given a context and an expression, however, our judgments synthesize an expression from a context and a type. The IREFINE judgments correspond to rules where a type would be checked and the EGUESS judgments to rules where a type would be synthesized.

3 Synthesis of Products

3.1 Overview of Products

In this section, I describe the process of adding k-ary (where k $>$ 1) tuples to the type-directed synthesis framework. The syntax of tuples is identical to that of OCaml. A tuple of k values is written:

$$(v_1, ..., v_k)$$

Tuples are eliminated using projection, which, in our rules, follows syntax similar to that of Standard ML. The j^{th} projection of a k-ary tuple (where $1 \leq j \leq k$)

$$\pi_j \ (v_1, ..., v_k)$$

evaluates to

$$v_j$$

As in OCaml, a tuple's type is written as the ordered cartesian product of the types of its constituents. For example, the type

$$\mathsf{nat} * \mathsf{list} * \mathsf{bool}$$

denotes the type of a tuple whose members are, in order, a natural number, a list, and a boolean.

3.2 Tuples as a Derived Form

It is entirely possible to encode tuples as a derived form in the original type-directed synthesis framework. We could declare an algebraic datatype with a single variant that stores k items of the necessary types. Projection would entail matching on a "tuple" to extract its contents. Although this strategy will successfully integrate a functional equivalent of products into the system, first class tuples are more desirable for several reasons.

Convenience. Since the existing system lacks polymorphism, we would need to declare a new algebraic type for every set of types we wish to combine into a product. Even with polymorphism, we would still need a new type for each arity we might desire. This approach would swiftly become tedious to a programmer and, worse, would lead the synthesizer to write incomprehensible programs filled with computer-generated type declarations and corresponding match statements.

Orthogonality. Implementing tuples as a derived form is an inelegant and syntactically-heavy way to add an essential language feature to the system. In contrast, first-class tuples are orthogonal, meaning we can add, remove, or modify other features without affecting products.

Cleaner theory. First class tuples simplify a number of other language features, including constructors and match statements. In the original type-directed synthesis system, variable-arity constructors were necessary to implement fundamental recursive types like lists. First class tuples can replace this functionality, streamlining the theory behind the entire system.

Efficient synthesis. Theorem-proving machinery can take advantage of the logical properties of product types to produce a more efficient synthesis algorithm (see Section 3.3). Although we could specially apply these techniques to single-variant, non-recursive algebraic datatypes (the equivalent of tuples), we can simplify the overall synthesis algorithm by delineating a distinct product type.

Informative matching. The original synthesis implementation only generates a match statement if doing so partitions the examples over two or more branches, a restriction known as *informative matching*. The rationale for this requirement is that pattern matching is only useful if it separates our synthesis task into smaller subproblems, each with fewer examples. A single-constructor variant would never satisfy this constraint, meaning that we would need another special case in order to use match statements to "project" on our tuple derived form.

τ	$::=$	$\mathcal{B} \mid \tau_1 \to \tau_2 \mid \mathsf{unit} \mid \tau_1 * ... * \tau_m$	(Types: base type, function type, unit type, and product type)
v	$::=$	$() \mid (v_1, ..., v_m)$ $\mid C\ v$ $\mid \mathsf{fix}\ f(x:\tau_1):\tau_2 = e$ $\mid pf$	(Values: unit, tuples, single-arity constructors, functions, and partial functions)
pf	$::=$	$\overline{v_i \Rightarrow \varepsilon_i}^{i \in m}$	(Partial functions: a finite map from inputs to outputs)
ε	$::=$	$() \mid (\varepsilon_1, ..., \varepsilon_m) \mid C\ \varepsilon \mid pf$	(Examples)
e	$::=$	$x \mid pf$ $\mid () \mid (e_1, ..., e_m) \mid \pi_k\ e$ $\mid \mathsf{fix}\ f(x:\tau_1):\tau_2 = e \mid e_1\ e_2$ $\mid C\ e$ $\mid \mathsf{match}\ e\ \mathsf{with}\ \overline{C_i\ x \to e_i}^{i \in l}$	(Expressions: projection is newly added)
E	$::=$	$x \mid E\ I \mid \pi_k\ E$	(Elimination forms: projection is newly added)
I	$::=$	E $\mid () \mid (I_1, ..., I_m)$ $\mid C\ I$ $\mid \mathsf{fix}\ f(x:\tau_1):\tau_2 = I$ $\mid \mathsf{match}\ E\ \mathsf{with}\ \overline{C_i\ x \to I_i}^{i \in l}$	(Introduction forms: unit and tuples are newly added)
Σ	$::=$	$\cdot \mid \Sigma, C : \tau \to \mathcal{B}$	(Constructor context: a map from constructor names to types)
$\Gamma, \mathcal{A}, \mathcal{F}$	$::=$	$\cdot \mid \Gamma, E : \tau$	(Type context: bindings of elimination forms to types)
$\mathcal{E}$	$::=$	$\cdot \mid \mathcal{E}, E \triangleright \varepsilon$	(Example context: elimination forms refined by examples)
w	$::=$	$\langle \mathcal{E}; \varepsilon \rangle$	(A single world: an example context and result example)
$\mathcal{W}$	$::=$	$\cdot \mid \mathcal{W}, w$	(A set of worlds)

Figure 5: Updated grammars for type-directed synthesis with products. Changes are highlighted.

$$\boxed{\Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Sigma_2 \mid \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2}$$

FOCUS-UNIT

$$\frac{}{\Sigma \mid \Gamma; \mathcal{A}; \mathcal{F}, E : \mathsf{unit} \mid \mathcal{W} \Longrightarrow \Sigma \mid \Gamma, E : \mathsf{unit}; \mathcal{A}; \mathcal{F} \mid \mathcal{W}}$$

FOCUS-BASE

$$\frac{}{\Sigma \mid \Gamma; \mathcal{A}; \mathcal{F}, E : \mathcal{B} \mid \mathcal{W} \Longrightarrow \Sigma \mid \Gamma, E : \mathcal{B}; \mathcal{A}; \mathcal{F} \mid \mathcal{W}}$$

FOCUS-FUN

$$\frac{}{\Sigma \mid \Gamma; \mathcal{A}; \mathcal{F}, E : \tau_1 \to \tau_2 \mid \mathcal{W} \Longrightarrow \Sigma \mid \Gamma, E : \tau_1 \to \tau_2; \mathcal{A}; \mathcal{F} \mid \mathcal{W}}$$

FOCUS-TUPLE

$$\frac{}{\begin{array}{l}\Sigma \mid \Gamma; \mathcal{A}; \mathcal{F}, E : \tau_1 * ... * \tau_m \mid \overline{\langle \mathcal{E}_i, E \triangleright (I_{(i,1)}, ..., I_{(i,m)}); \varepsilon_i \rangle}^{i \in n} \Longrightarrow \\ \Sigma \mid \Gamma; \mathcal{A}, E : \tau_1 * ... * \tau_m; \mathcal{F}, \overline{\pi_j\ E : \tau_j}^{j \in m} \mid \overline{\langle \mathcal{E}_i, \overline{\pi_j\ E \triangleright I_{(i,j)}}^{j \in m}; \varepsilon_i \rangle}^{i \in n}\end{array}}$$

Figure 6: Focusing rules for type-directed synthesis with products.

$$\boxed{\Sigma \mid \Gamma; \mathcal{A}; \mathcal{F} \mid \mathcal{W} \vdash \tau \overset{I}{\rightsquigarrow} I}$$

IREFINE-GUESS
$$\frac{\Sigma \mid \Gamma; \mathcal{A}; \cdot \vdash \mathcal{B} \overset{E}{\rightsquigarrow} E \qquad \forall i \in n,\ \mathcal{E}_i(E) \to^* \varepsilon_i}{\Sigma \mid \Gamma; \mathcal{A}; \cdot \mid \overline{\langle \mathcal{E}_i; \varepsilon_i \rangle}^{i \in n} \vdash \mathcal{B} \overset{I}{\rightsquigarrow} E}$$

IREFINE-UNIT
$$\frac{}{\Sigma \mid \Gamma; \mathcal{A}; \cdot \mid \overline{\langle \mathcal{E}_i; () \rangle}^{i \in n} \vdash \mathsf{unit} \overset{I}{\rightsquigarrow} ()}$$

IREFINE-FIX
$$\frac{\Sigma \mid \Gamma; \mathcal{A}; f : \tau_1 \to \tau_2, x : \tau_1 \mid \overline{\langle \mathcal{E}_i, f \triangleright \overline{v_{(i,j)} \Rightarrow \varepsilon_{(i,j)}}^{j \in m_i}, x \triangleright v_{(i,k)}; \varepsilon_{(i,k)} \rangle}^{(i,k) \in (n, m_i)} \vdash \tau_2 \overset{I}{\rightsquigarrow} I}{\Sigma \mid \Gamma; \mathcal{A}; \cdot \mid \overline{\langle \mathcal{E}_i; \overline{v_{(i,j)} \Rightarrow \varepsilon_{(i,j)}}^{j \in m_i} \rangle}^{i \in n} \vdash \tau_1 \to \tau_2 \overset{I}{\rightsquigarrow} \mathsf{fix}\ f(x : \tau_1) : \tau_2 = I}$$

IREFINE-CTOR
$$\frac{C : \tau \to \mathcal{B} \in \Sigma \qquad \Sigma \mid \Gamma; \mathcal{A}; \cdot \mid \overline{\langle \mathcal{E}_i; \varepsilon_i \rangle}^{i \in n} \vdash \tau \overset{I}{\rightsquigarrow} I}{\Sigma \mid \Gamma; \mathcal{A}; \cdot \mid \overline{\langle \mathcal{E}_i; C\ \varepsilon_i \rangle}^{i \in n} \vdash \mathcal{B} \overset{I}{\rightsquigarrow} C\ I}$$

IREFINE-TUPLE
$$\frac{\forall i \in m,\ \Sigma \mid \Gamma; \mathcal{A}; \cdot \mid \overline{\langle \mathcal{E}_j; \varepsilon_{(i,j)} \rangle}^{j \in n} \vdash \tau_i \overset{I}{\rightsquigarrow} I_i}{\Sigma \mid \Gamma; \mathcal{A}; \cdot \mid \overline{\langle \mathcal{E}_j; (\varepsilon_{(1,j)}, \ldots, \varepsilon_{(m,j)}) \rangle}^{j \in n} \vdash \tau_1 * \ldots * \tau_m \overset{I}{\rightsquigarrow} (I_1, \ldots, I_m)}$$

IREFINE-MATCH
$$\frac{\begin{array}{c}\Sigma \mid \Gamma; \mathcal{A}; \cdot \vdash \mathcal{B} \overset{E}{\rightsquigarrow} E \\ \forall j \in m,\ C_j : \tau_j \to \mathcal{B} \in \Sigma \qquad \forall j \in m,\ \Sigma \mid \Gamma; \mathcal{A}; x : \tau_j \mid \{ \langle \mathcal{E}_i, x \triangleright v; \varepsilon_i \rangle \mid \mathcal{E}_i(E) \to^* C_j\ v \} \vdash \tau \overset{I}{\rightsquigarrow} I_j\end{array}}{\Sigma \mid \Gamma; \mathcal{A}; \cdot \mid \overline{\langle \mathcal{E}_i; \varepsilon_i \rangle}^{i \in n} \vdash \tau \overset{I}{\rightsquigarrow} \mathsf{match}\ E\ \mathsf{with}\ \overline{C_j\ x \to I_j}^{j \in m}}$$

IREFINE-FOCUS
$$\frac{\Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Sigma_2 \mid \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2 \qquad \Sigma_2 \mid \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2 \vdash \tau \overset{I}{\rightsquigarrow} I}{\Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \vdash \tau \overset{I}{\rightsquigarrow} I}$$

$$\boxed{\Sigma \mid \Gamma; \mathcal{A}; \mathcal{F} \vdash \tau \overset{E}{\rightsquigarrow} E}$$

EGUESS-CTX
$$\frac{E : \mathcal{B} \in \Gamma}{\Sigma \mid \Gamma; \mathcal{A}; \cdot \vdash \mathcal{B} \overset{E}{\rightsquigarrow} E}$$

EGUESS-APP
$$\frac{\Sigma \mid \Gamma, E_1 : \tau_1 \to \tau_2; \mathcal{A}; \cdot \mid \cdot \vdash \tau_1 \overset{I}{\rightsquigarrow} I_1 \qquad \Sigma \mid \Gamma, E_1 : \tau_1 \to \tau_2; \mathcal{A}; E_1\ I_1 : \tau_2 \vdash \tau \overset{E}{\rightsquigarrow} E}{\Sigma \mid \Gamma, E_1 : \tau_1 \to \tau_2; \mathcal{A}; \cdot \vdash \tau \overset{E}{\rightsquigarrow} E}$$

EGUESS-FOCUS
$$\frac{\Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \cdot \Longrightarrow \Sigma_2 \mid \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \cdot \qquad \Sigma_2 \mid \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \vdash \tau \overset{E}{\rightsquigarrow} E}{\Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \vdash \tau \overset{E}{\rightsquigarrow} E}$$

Figure 7: The updated judgments for type-directed synthesis with products. Altered judgments have been highlighted.

3.3 Focusing

Theorem proving strategy. Product types are the equivalents of logical conjunction, which itself is *left-invertible* [15]. That is, when we know that $A \wedge B$ is true, we can eagerly conclude that A and B are individually true without loss of information. Should we later need the fact that $A \wedge B$ is true, we can easily re-prove it.

We can turn this invertibility into an efficient strategy for theorem-proving called *focusing*. When we add a value of product type to the context (our proof-term equivalent of statements we know to be true) we can eagerly project on it in order to break it into its non-product type constituents. Since products are invertible, we can always reconstruct the original entity from its projected components if the need arises. Where many parts of the type-directed synthesis strategy require us to guess derivations that might not lead to a candidate program, focusing allows us to advance our proof search in a way that will never need to be backtracked.

Synthesis strategy. Integrating focusing into type-directed synthesis requires several fundamental changes to the existing theory, including the addition of an entirely new focusing judgment form (Figure 6) and additional "focusing" contexts in the IREFINE and EGUESS judgments (Figure 7). The grammar of the type context Γ has subtly changed: rather than binding names to types, it now binds elimination forms E (Figure 5). This alteration allows the context to store, not only a value $(x : \tau_1 * \tau_2)$, but also the results of focusing: $(\pi_1\ x : \tau_1)$ and $(\pi_2\ x : \tau_2)$.

Two new contexts share the same structure as Γ: a *focusing context* $\mathcal{F}$ and an *auxiliary context* $\mathcal{A}$. We use the contexts as follows: when a new name or, more generally, elimination form is to be added to the context, we first insert it into $\mathcal{F}$. Items in $\mathcal{F}$ are repeatedly focused and reinserted into $\mathcal{F}$ until they cannot be focused any further, at which point they move into Γ. The terms that we project upon in the course of focusing are no longer useful for the synthesis process but are critical for proving many properties about our context; they are therefore permanently moved into $\mathcal{A}$.

Invariants. We can crystalize this workflow into several invariants about the contexts:

1. New names and elimination forms are added to $\mathcal{F}$.
2. Elimination forms in Γ cannot be focused further.
3. No other rule can be applied until focusing is done.
4. Only elimination forms in Γ can be used to synthesize expressions.

We have already discussed the justification for the first two invariants. The third invariant ensures that we prioritize backtracking-free focusing over the nondeterministic choices inherent in applying other rules. By always focusing first, we add determinism to our algorithm and advance our proof search as far as possible before making any decisions that we might later have to backtrack. To enforce this invariant, all of the IREFINE and EGUESS rules that synthesize expressions require $\mathcal{F}$ to be empty.

The final invariant ensures that we never generate terms from $\mathcal{A}$ or $\mathcal{F}$ that may not have been entirely focused, preserving the property that we only synthesize normal forms. Otherwise, we could synthesize both a term $(x : \tau_1 * \tau_2)$ and the eta-equivalent expression $((\pi_1\ x, \pi_2\ x) : \tau_1 * \tau_2)$. This invariant implies an eta-long normal form for tuples: among the aforementioned two expressions, our algorithm will always synthesize the latter.

Focusing judgments. Special IREFINE-FOCUS and EGUESS-FOCUS rules allow the focusing process to advance to completion so that synthesis can continue. They make calls to a separate focusing judgment (Figure 6) that operates on contexts and example worlds. The judgment

$$\Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Sigma_2 \mid \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2$$

states that

> *The contexts Σ_1, Γ_1, $\mathcal{A}_1$, and $\mathcal{F}_1$ along with example worlds $\mathcal{W}_1$ can be focused into the contexts Σ_2, Γ_2, $\mathcal{A}_2$, and $\mathcal{F}_2$ along with example worlds $\mathcal{W}_2$.*

The focusing rules move any elimination form of non-product type from $\mathcal{F}$ to Γ without modifying any other contexts. On product types, the rule FOCUS-TUPLE projects on an elimination form in $\mathcal{F}$, returning the resulting expressions to $\mathcal{F}$ for possible additional focusing. It splits the examples of the focused expression amongst the newly-created projections and moves the now-unnecessary original expression into $\mathcal{A}$.

3.4 Changes to Synthesis Judgments

The updated grammars, focusing rules, and judgments for type-directed synthesis with products appear in Figures 5, 6, and 7 respectively. Below are summaries of the changes that were made to the original theory in the process of adding products.

Tuple and projection expressions. Expressions for tuples and projection have been added to the grammar of expressions (e). Likewise, the product type has been added to the grammar of types (τ). Tuple creation is an introduction form with projection as the corresponding elimination form.

Single-argument constructors. Now that tuples exist as a separate syntactic construct, the variable-arity constructors of the original type-directed synthesis judgments are redundant: creating a constructor with an argument of product type achieves the same effect. Without products, variable-arity constructors were essential for defining recursive types that stored values (i.e., lists), but this functionality can now entirely be implemented using tuples. Therefore, in the updated judgments, all constructors have a single argument, enormously simplifying the IREFINE-CTOR and IREFINE-MATCH rules.

One drawback of this choice is that pattern matching can no longer deconstruct both a constructor and the implicit tuple embedded inside it as in the old judgments. Instead, pattern matching always produces a single variable that must be focused and projected upon in order to extract its components. Additional pattern-matching machinery could be added to the theory (and is present in the implementation) to resolve this shortcoming.

Unit type. Now that all constructors require a single argument, what is the base case in an inductive datastructure? Previously, constructors with no arguments (`O`, `Nil`, etc.) served this purpose, but the new rules require them to have an argument as well. To resolve this problem, a unit type (`unit`) and unit value (written `()`) have been introduced. Former no-argument constructors now have type `unit`. A straightforward IREFINE-UNIT rule synthesizes the unit value at unit type.

Tuple and projection judgments. A new IREFINE-TUPLE judgment synthesizes a tuple at product type when all examples are tuple values with the same arity, inheriting much of the behavior previously included in the IREFINE-CTOR rule. It creates one subproblem for every value contained within the tuple and partitions the examples accordingly. Notably, there is no corresponding EGUESS-PROJ judgment. All projection instead occurs during focusing.

Corrected application judgment. The absence of an EGUESS-PROJ rule required changes to the EGUESS-APP judgment. As a motivating example, consider the list *unzip* function:

```
let unzip (ls:pairlist) :  list * list =
   match ls with
   | [] -> ([], [])
   | (a, b) ::  tl -> (a ::  #1 (unzip tl), b ::  #2 (unzip tl))
```

Synthesizing *unzip* requires projecting on a function application (highlighted above).

The application rule in the original type-directed synthesis judgments is reproduced below. It has been modified slightly to include the new context structure accompanying the updated judgments in Figure 7:

$$\frac{\Sigma \mid \Gamma; \mathcal{A}; \cdot \vdash \tau_1 \to \tau \overset{E}{\rightsquigarrow} E \qquad \Sigma \mid \Gamma; \mathcal{A}; \cdot \mid \cdot \vdash \tau_1 \overset{I}{\rightsquigarrow} I}{\Sigma \mid \Gamma; \mathcal{A}; \cdot \vdash \tau \overset{E}{\rightsquigarrow} E\ I}$$

This judgment requires that we synthesize applications at the goal type and immediately use them to solve the synthesis problem. In *unzip*, however, we need to synthesize an application at a product type and subsequently project on it, meaning that, with this version of the EGUESS-APP judgment, we cannot synthesize *unzip*.

Alternatively, we could add a straightforward EGUESS-PROJ judgment that would allow us to project on applications:

$$\frac{\Sigma \mid \Gamma; \mathcal{A}; \cdot \vdash \tau_1 * \ldots * \tau_m \overset{E}{\rightsquigarrow} E \qquad 1 \leq k \leq m}{\Sigma \mid \Gamma; \mathcal{A}; \cdot \vdash \tau_k \overset{E}{\rightsquigarrow} \pi_k\ E}$$

Doing so, however, undermines the invariant that we only generate normal forms. Suppose we wished to synthesize an expression at type $\tau_1 * \tau_2$. We could either synthesize an application $E\ I : \tau_1 * \tau_2$ or the corresponding eta-long version $(\pi_1\ (E\ I), \pi_2\ (E\ I)) : \tau_1 * \tau_2$.

For a solution to this quandary, we need to return to the underlying proof theory. In the sequent calculus, the left rule for implication, which corresponds to our EGUESS-APP judgment, appears as below [6, 15]:

$$\frac{(1)\ \Gamma, A \to B \Longrightarrow A \qquad (2)\ \Gamma, A \to B, B \Longrightarrow C}{(3)\ \Gamma, A \to B \Longrightarrow C}$$

That is:

> *If the context contains proof that A implies B and (1) we can prove A and (2) if the context also contains a proof of B, then we can prove C, then (3) we can prove C.*

This rule suggests a slightly different application rule, namely the one in Figure 7:

EGUESS-APP

$$\frac{\Sigma \mid \Gamma, E_1 : \tau_1 \to \tau_2; \mathcal{A}; \cdot \mid \cdot \vdash \tau_1 \overset{I}{\rightsquigarrow} I_1 \qquad \Sigma \mid \Gamma, E_1 : \tau_1 \to \tau_2; \mathcal{A}; E_1\ I_1 : \tau_2 \vdash \tau \overset{E}{\rightsquigarrow} E}{\Sigma \mid \Gamma, E_1 : \tau_1 \to \tau_2; \mathcal{A}; \cdot \vdash \tau \overset{E}{\rightsquigarrow} E}$$

The new EGUESS-APP rule allows us to generate an application at any type, not just the goal type, and to add it to the context. We then focus the application, thereby projecting on it as necessary.

$$\textsc{Focus-Closure-Base}\quad \frac{}{\Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow^* \Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1}$$

$$\textsc{Focus-Closure-Step}\quad \frac{\Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Sigma_2 \mid \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2 \qquad \Sigma_2 \mid \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2 \Longrightarrow^* \Sigma_3 \mid \Gamma_3; \mathcal{A}_3; \mathcal{F}_3 \mid \mathcal{W}_3}{\Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow^* \Sigma_3 \mid \Gamma_3; \mathcal{A}_3; \mathcal{F}_3 \mid \mathcal{W}_3}$$

Figure 8: Judgments for the transitive closure of focusing.

$$\textsc{T-Var}\quad \frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau} \qquad \textsc{T-Abs}\quad \frac{\Gamma, f : \tau_1 \to \tau_2, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash \mathsf{fix}\ f(x : \tau_1) : \tau_2 = e : \tau_1 \to \tau_2} \qquad \textsc{T-App}\quad \frac{\Gamma \vdash e_1 : \tau_1 \to \tau_2 \qquad \Gamma \vdash e_2 : \tau_1}{\Gamma \vdash e_1\ e_2 : \tau_2}$$

$$\textsc{T-Unit}\quad \frac{}{\Gamma \vdash () : \mathsf{unit}} \qquad \textsc{T-Tuple}\quad \frac{\forall\, i \in m,\ \Gamma \vdash e_i : \tau_i}{\Gamma \vdash (e_1, ..., e_m) : \tau_1 * ... * \tau_m} \qquad \textsc{T-Proj}\quad \frac{\Gamma \vdash e : \tau_1 * ... * \tau_m \qquad 1 \leq k \leq m}{\Gamma \vdash \pi_k\ e : \tau_k}$$

$$\textsc{T-Ctor}\quad \frac{C : \tau \to \mathcal{B} \in \Sigma \qquad \Gamma \vdash e : \tau}{\Gamma \vdash C\ e : \mathcal{B}} \qquad \textsc{T-Match}\quad \frac{\forall\, i \in m,\ C_i : \tau_i \to \mathcal{B} \in \Sigma \qquad \Gamma \vdash e : \mathcal{B} \qquad \forall\, i \in m,\ \Gamma, x : \tau_i \vdash e_i : \tau}{\Gamma \vdash \mathsf{match}\ e\ \mathsf{with}\ \overline{C_i\ x \to e_i}^{\,i \in m} : \tau}$$

Figure 9: Typing judgments for the simply typed lambda calculus with recursion, algebraic datatypes, and products.

This application now becomes available for use in the synthesis problem on which we are currently working. In effect, we take a forward step in the context by guessing an application.

Note that, although the new application judgment corresponds more closely to the proof theory, the application judgment in the original type-directed synthesis rules was not incorrect. Without projection, the only possible use for a function application in the context would be to immediately use it to satisfy the synthesis goal. The application judgment therefore simply short-circuited the process of first adding an application to the context and then using it via EGuess-Ctx.

3.5 Properties

Overview

This section describes a number of useful properties of the type-directed synthesis system with products and focusing. The accompanying proofs and lemmas appear in Appendix A. Note that, since none of the synthesis or focusing judgments modify the constructor context (Σ), it has been elided from the following theorems for readability; for our purposes, it is assumed to be a fixed and globally available entity.

Theorem 2.2: Progress of Focusing

$$\mathcal{F}_1 = \cdot \quad \text{OR} \quad \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2$$

Whenever the focusing context is non-empty, a focusing judgment can be applied.

$$\text{S-CTOR}\ \frac{e \to e'}{C\ e \to C\ e'} \qquad \text{S-TUPLE}\ \frac{e_k \to e'_k}{(v_1,...,e_k,...,e_n) \to (v_1,...,e'_k,...,e_n)} \qquad \text{S-PROJ1}\ \frac{e \to e'}{\pi_k\ e \to \pi_k\ e'} \qquad \text{S-PROJ2}\ \frac{1 \le k \le m}{\pi_k\ (e_1,\ ...,\ e_m) \to e_k}$$

$$\text{S-APP1}\ \frac{e_1 \to e'_1}{e_1\ e_2 \to e'_1\ e_2} \qquad \text{S-APP2}\ \frac{}{(\mathsf{fix}\ f(x:\tau_1):\tau_2 = e_1)\ e_2 \to e_1[x \mapsto e_2, f \mapsto \mathsf{fix}\ f(x:\tau_1):\tau_2 = e_1]}$$

$$\text{S-MATCH1}\ \frac{e \to e'}{\mathsf{match}\ e\ \mathsf{with}\ \overline{C_i\ x \to e_i}^{i \in m} \to \mathsf{match}\ e'\ \mathsf{with}\ \overline{C_i\ x \to e_i}^{i \in m}} \qquad \text{S-MATCH2}\ \frac{1 \le k \le m}{\mathsf{match}\ C_k\ v\ \mathsf{with}\ \overline{C_i\ x \to e_i}^{i \in m} \to e_k[x \mapsto v]}$$

Figure 10: Small-step semantics for the simply typed lambda calculus with recursion, algebraic datatypes, and products.

$$\text{S-CLOSURE-BASE}\ \frac{}{e \to^* e} \qquad \text{S-CLOSURE-STEP}\ \frac{e \to e' \qquad e' \to^* e''}{e \to^* e''}$$

Figure 11: Judgments for the transitive closure of the small step semantics.

$$\text{TYPE-CTX-EMPTY-WF}\ \frac{}{\vdash \cdot\ \ wf} \qquad \text{TYPE-CTX-ONE-WF}\ \frac{\vdash \Gamma\ \ wf \qquad \vdash E : \tau}{\vdash \Gamma, E : \tau\ \ wf}$$

Figure 12: Judgments for the well-formedness of a type or focusing context.

$$\boxed{[\Gamma]_{vars}}$$

$$[\cdot]_{vars} = \{\} \qquad [\Gamma, x:\tau]_{vars} = \{x:\tau\} \cup [\Gamma]_{vars} \qquad [\Gamma, e:\tau]_{vars} = [\Gamma]_{vars}\ \textit{where}\ e \neq x$$

Figure 13: Rules for the variable extraction relation, which produces the set of all type bindings of variables from a context that might otherwise contain arbitrary elimination forms.

Theorem 2.3: Preservation of Focusing

$$\begin{array}{ll}\textsc{If} & \vdash \Gamma_1 \ \ wf, \quad \vdash \mathcal{A}_1 \ \ wf, \quad \vdash \mathcal{F}_1 \ \ wf \\ \textsc{and} & \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2 \\ \textsc{then} & \vdash \Gamma_2 \ \ wf, \quad \vdash \mathcal{A}_2 \ \ wf, \quad \vdash \mathcal{F}_2 \ \ wf \end{array}$$

The application of a focusing judgment preserves the well-formedness of Γ, $\mathcal{A}$, and $\mathcal{F}$. The judgments for well-formedness are in Figure 12.

Theorem 2.5: Termination of Focusing

$$\Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow^* \Gamma_2; \mathcal{A}_2; \cdot \mid \mathcal{W}_2$$

The focusing process always terminates. The judgments for the transitive closure of focusing are in Figure 8.

Theorem 3.1: Focusing is Deterministic

$$\begin{array}{lllll}\textsc{If} & \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow^* \Gamma_{2a}; \mathcal{A}_{2a}; \cdot \mid \mathcal{W}_2 \\ \textsc{and} & \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow^* \Gamma_{2b}; \mathcal{A}_{2b}; \cdot \mid \mathcal{W}_2 \\ \textsc{then} & \Gamma_{2a} = \Gamma_{2b} & \textsc{and} & \mathcal{A}_{2a} = \mathcal{A}_{2b} \end{array}$$

The focusing process proceeds deterministically.

Theorem 3.2: Order of Focusing

Define the judgments $\Gamma; \mathcal{A}; \mathcal{F} \vdash \tau \overset{E'}{\rightsquigarrow} E$ and $\Gamma; \mathcal{A}; \mathcal{F} \mid \mathcal{W} \vdash \tau \overset{I'}{\rightsquigarrow} I$ to be identical to those in Figure 7 except that $\mathcal{F}$ need not be empty for any synthesis rule to be applied.

$$\begin{array}{lll} \Gamma; \mathcal{A}; \mathcal{F} \mid \mathcal{W} \vdash \tau \overset{I}{\rightsquigarrow} I & \textsc{iff} & \Gamma; \mathcal{A}; \mathcal{F} \mid \mathcal{W} \vdash \tau \overset{I'}{\rightsquigarrow} I \\ \Gamma; \mathcal{A}; \mathcal{F} \vdash \tau \overset{E}{\rightsquigarrow} E & \textsc{iff} & \Gamma; \mathcal{A}; \mathcal{F} \vdash \tau \overset{E'}{\rightsquigarrow} E \end{array}$$

If we make the proof search process more non-deterministic by permitting arbitrary interleavings of focusing and synthesis, we do not affect the set of expressions that can be generated. This more lenient process makes many of the other theorems simpler to prove.

Admissibility of Focusing

Define the judgments $\Gamma \vdash \tau \overset{E}{\rightsquigarrow} E$ and $\Gamma \mid \mathcal{W} \vdash \tau \overset{I}{\rightsquigarrow} I$ to be identical to those in Figure 7 without focusing. Instead of focusing to project on tuples, we introduce the EGuess-Proj rule as below:

$$\frac{\Gamma, E_1 : \tau_1 * \ldots * \tau_m,\ \pi_1\ E_1 : \tau_1,\ \ldots,\ \pi_m\ E_1 : \tau_m \vdash \tau \overset{E}{\rightsquigarrow} E}{\Gamma, E_1 : \tau_1 * \ldots * \tau_m \vdash \tau \overset{E}{\rightsquigarrow} E}\ \textsc{EGuess-Proj}$$

Theorem 4.1: Soundness of Focusing

IF $\Gamma;\cdot;\cdot \mid \cdot \vdash \tau \overset{I}{\rightsquigarrow} I$ THEN $\Gamma \mid \cdot \vdash \tau \overset{I}{\rightsquigarrow} I$

IF $\Gamma;\cdot;\cdot \vdash \tau \overset{E}{\rightsquigarrow} I$ THEN $\Gamma \vdash \tau \overset{E}{\rightsquigarrow} I$

Any judgment that can be proven in the system with focusing can be proven in the system without focusing. That is, the system with focusing is no more powerful than the system without focusing.

Theorem 4.2: Completeness of Focusing

IF $\Gamma \mid \cdot \vdash \tau \overset{I}{\rightsquigarrow} I$ THEN $\Gamma;\cdot;\cdot \mid \cdot \vdash \tau \overset{I}{\rightsquigarrow} I$

IF $\Gamma \vdash \tau \overset{E}{\rightsquigarrow} I$ THEN $\Gamma;\cdot;\cdot \vdash \tau \overset{E}{\rightsquigarrow} I$

Any judgment that can be proven in the system without focusing can be proven in the system with focusing. That is, the system with focusing is at least as powerful as the system without focusing. Together, Theorems 4.1 and 4.2 demonstrate that the system with focusing is exactly as powerful as the system without focusing

Note 4.3: Pushing Examples Through Elimination Forms

Observe that we can push examples through the tuple focusing judgments but not the EGUESS-PROJ judgment. The sole distinction between these two judgments is where they may be used in the synthesis process: EGUESS-PROJ must be used during the elimination-guessing phase while focusing may take place at any point in the synthesis process. This discrepancy demonstrates that it is possible, at least in some cases, to make use of example information when synthesizing elimination forms. Since guessing elimination forms requires raw term enumeration, any constraints on this process would yield significant performance improvements. I leave further investigation of this behavior to future work.

Theorem 5.1: Soundness of Synthesis Judgments

IF $\Gamma;\mathcal{A};\mathcal{F} \mid \overline{\langle \mathcal{E};\varepsilon \rangle}^{n} \vdash \tau \overset{I}{\rightsquigarrow} I$

THEN $[\Gamma,\mathcal{A},\mathcal{F}]_{vars} \vdash I : \tau$ AND $\forall\, i \in n,\ \mathcal{E}_i(I) \rightarrow^* \varepsilon_i$

IF $\Gamma;\mathcal{A};\mathcal{F} \vdash \tau \overset{E}{\rightsquigarrow} E$

THEN $[\Gamma,\mathcal{A},\mathcal{F}]_{vars} \vdash E : \tau$

The type-directed synthesis system will always produce expressions with the desired type that obey the input-output examples. The typing judgments, small-step semantics, and variable extraction relation are defined in Figures 9, 10, and 13.

Completeness of Synthesis Judgments

Stating a completeness theorem for type-directed synthesis, let alone proving it, remains an open question. I therefore do not attempt to prove completeness here.

Program	Examples	Time (s)
`make_pair`	2	0.001
`make_triple`	2	0.003
`make_quadruple`	2	0.003
`fst`	2	0.003
`snd`	2	0.003
`unzip (bool)`	7	0.005
`unzip (nat)`	7	0.006
`zip (bool)`	28	1.183
`zip (nat)`	28	2.879

(a) The number of examples and seconds necessary to synthesize various canonical programs that make use of tuples and projection.

Program	Before (s)	After (s)
`bool_band`	0.006	0.007
`list_append`	0.016	0.024
`list_fold`	0.304	0.183
`list_filter`	2.943	0.657
`list_length`	0.003	0.003
`list_map`	0.001	0.039
`list_nth`	0.061	0.072
`list_sorted_insert`	91.427	10.600
`list_take`	0.197	0.083
`nat_add`	0.016	0.011
`nat_max`	0.037	0.065
`nat_iseven`	0.005	0.006
`tree_binsert`	0.762	0.865
`tree_nodes_at_level`	1.932	0.812
`tree_preorder`	0.081	0.076

(b) A comparison of the synthesis performance between the system as described in the original type-directed synthesis paper (Before) and the updated version of the system that includes the changes described in this thesis (After).

Figure 14: Performance evaluation of the prototype implementation.

4 Evaluation

4.1 Implementation Changes

The prototype implementation of the type-directed synthesis system consists of approximately 3,000 lines of OCaml. The structure of the implementation largely mirrors that of the synthesis judgments, with type and constructor contexts and recursive proof search for a program that satisfies the input-output examples. The implementation deviates only when searching for elimination forms, during which it constructs terms bottom-up instead of the less-efficient top-down guessing suggested by the synthesis judgments.

To add tuples, I altered between 1,500 and 2,000 lines of code. Notable changes included:

- Modifying the parser and type-checker to add tuples, projection, and single-argument constructors.
- Entirely rewriting the type context to store elimination forms and perform focusing.
- Updating the synthesis process to conform to the new judgments.
- Adding a unit type and value throughout the system.
- Extending the pattern-matching language to include mechanisms for recursively matching on the tuples embedded within constructors.
- Refactoring, simplification, modularization, dead code elimination, and other cleanup to prepare the system for future extensions.

4.2 Performance Evaluation

Raw Performance. Figure 14a summarizes the performance of the prototype implementation on a number of canonical programs that make use of tuples and projection. Programs that did not require recursive calls, like `make_pair` and `fst`, were synthesized instantly. The list `unzip` function was likewise generated in a small fraction of a second, but required a relatively large number of examples. The synthesizer would overfit if any example was discarded. The list `zip` function took the longest to synthesize and required far more examples. The reason for this behavior is that `zip` must handle many more cases than `unzip`, such as the possibility that the two lists are of different lengths.

The `unzip` and `zip` programs were tested twice: once each with lists comprising boolean and natural number values. The natural number case predictably took longer, since the synthesizer likely spent a portion of its time generating programs that manipulate the values stored in the lists as opposed to the lists themselves. Since there are a finite number of boolean values, the search space of these faulty programs was much smaller. With polymorphic lists, this behavior could be eliminated completely.

Performance Regression. The changes made in the process of adding tuples affected the entire synthesis system. Even programs that do not directly make use of tuples could take longer to synthesize. All multi-argument constructors, including `Cons`, now require tuples when created and projection when destroyed. No-argument constructors, which formerly served as the base cases of the synthesis process, now store values of unit type. Whenever a variable is added to the context, it goes through the process of focusing before it can be used. While small, these modifications add overhead that impacts any program we might attempt to synthesize.

Figure 14b contains a performance comparison of the system as described in the original paper with the updated version that includes tuples. The running times were generated using the same computer on the same date; they are not lifted directly from the type-directed synthesis paper [14]. The selected programs are a subset of those tested in the original paper. Overall, it appears that there has been little tangible change in the performance of the system. Most examples saw at most slight fluctuations. Several of the longer-running examples, however, experienced dramatic performance improvements, which might have resulted from changes made to the memoization behavior of the system when rewriting the type context.

5 Related Work

5.1 Solver-free Synthesis

StreamBit. Among the first modern synthesis techniques was sketching [19], developed by Armando Solar-Lezama, which allows users to specify an incomplete program (a *sketch*) that is later filled in by a synthesizer. Although this line of work evolved into a general-purpose, solver-aided synthesis language (see Section 5.2), the initial prototype was designed for programs that manipulate streams of bits (*bitstreaming* programs), a notable example of which is a cryptographic cipher. In the system's language, StreamBit, users specify composable transformers, splits, and joins on streams of bits that together form a program. The paper states that StreamBit helped programmers design optimized ciphers that were several times faster than existing versions written in C.

The process of synthesis in StreamBit involves two steps. First, a user writes a correctness specification in the form of a program that implements the desired functionality. A user then writes a sketch that describes some, but not all, of the program to be synthesized. The remaining integer and boolean holes are filled by the synthesizer (using mathematical transformations and linear algebra) to create a program whose behavior is equivalent to the specification. The theory behind this design is that it is easy to write the naive, unoptimized versions of bitstreaming programs that serve as specifications but difficult to manage the low-level intricacies of optimized variants. Sketching allows a human to describe the algorithmic insights behind possible optimizations while leaving the synthesizer to perform the complex task of working out the details.

Sketching also permits a programmer to determine that a particular syntactic template is *not* a viable solution should the synthesizer fail to generate a program. (The idea of discovering an algorithm in top-down fashion by testing syntactic templates for viability with the aid of a synthesizer or - more broadly - a solver is discussed at length in [4].)

Flash Fill. Designed by Sumit Gulwani, Flash Fill [7] synthesizes string transformations using input-output examples. The system is featured in Excel 2013, where it generalizes user-completed data transformations on a small set of examples into a formula that works for the larger spreadsheet. The synthesizer behind Flash Fill relies on a domain-specific language with constructs for substring manipulation, regular expression matching, loops, and conditionals.

Flash Fill works by examining each example in turn. It uses a graph to store the set of all possible transformations that could convert each input into its corresponding output. By taking the intersection of these graphs among all examples, Flash Fill synthesizes a transformation that works in the general case. This process relies on a variety of string manipulation algorithms and relations between transformations rather than a solver.

Flash Fill epitomizes several key synthesis design concepts. Like type-directed synthesis, Flash Fill derives a program from input-output examples, which are more convenient for programmers and non-programmers alike, rather than a complete specification. It uses a compact representation - namely a transformation graph - to capture a large number of possible programs in a small amount of storage. Finally, it synthesizes programs over a specialized, domain-specific language whose search space is smaller than that of a general-purpose language, reducing the difficulty of the overall synthesis task.

Escher. The closest relative to type-directed synthesis is Gulwani's general-purpose synthesis system, Escher [2]. Like type-directed synthesis, Escher aims to generate general-purpose, recursive programs from input-output examples, although Escher assumes the existence of an oracle that can be queried when the synthesizer needs additional guidance. Gulwani notes that Escher's synthesis process is "generic" in that it works over an arbitrary domain of instructions, although he only demonstrates the system on a language with algebraic datatypes.

In contrast to type-directed synthesis, Escher uses bottom-up *forward search*, generating programs of increasing size out of instructions, recursive calls to the function being synthesized, input variables, and other language atoms. Escher maintains a *goal graph*, which initially comprises a single node that stores the set of outputs a synthesized program would need to produce to satisfy the examples and solve the synthesis problem. When a generated program p satisfies some (but not all) of the outputs in a goal node, Escher creates a conditional that partitions the program's control flow, directing the satisfied cases to p. Doing so creates two new goal nodes: one for the condi-

tional's guard and the other to handle the cases that p failed to satisfy. When no more unsatisfied goal nodes remain, Escher has synthesized a program.

5.2 Solver-aided Synthesis

Syntax-guided Synthesis. Published in 2013 and co-authored by many leading researchers in solver-aided synthesis (including Rastislav Bodik, Sanjit Seshia, Rishabh Singh, Armando Solar-Lezama, Emina Torlak, and others), *Syntax-Guided Synthesis* (SyGuS) [3] represents an attempt to codify a "unified theory" of synthesis. The authors seek to create a standardized input format for synthesis queries, similar to those developed for SMT and SAT solvers. Doing so for SMT and SAT solvers has led to shared suites of benchmarks, objective performance comparisons between solvers, and annual solver competitions, an effort credited with encouraging researchers to advance the state of the art. More importantly, a common interface means that higher-level applications can call off-the-shelf solvers, abstracting away solvers as subroutines. When solver performance improves or better solvers are developed, applications instantly benefit. The SyGuS project seeks to create a similar environment for synthesizers [1].

The common synthesis format encodes three pieces of information: a *background theory* in which the synthesis problem will be expressed, a *correctness specification* dictating the requirements of an acceptable solution, and the *grammar* from which candidate programs may be drawn.

The paper includes a thorough taxonomy of modern synthesis techniques. It distinguishes between *deductive synthesis*, in which a synthesizer creates a program by proving its logical specification, and *inductive synthesis*, which derives a program from input-output examples. It also describes a popular inductive synthesis algorithm, *counterexample-guided inductive synthesis* (CEGIS). In CEGIS, a solver repeatedly generates programs that it attempts to verify against a specification. Whenever verification fails, the synthesis procedure can use the verifier's counterexample as an additional constraint in the program-generation procedure. If no new program can be constructed, synthesis fails; if no counterexample can be found, synthesis succeeds. In practice, most synthesis problems require only a few CEGIS steps to generate the counterexamples necessary to discover a solution.

Finally, the authors present an initial set of benchmarks in the common synthesis format and a performance comparison of existing synthesis systems.

In the sections that follow, I detail several lines of research that fit into the framework of solver-aided syntax-guided synthesis.

Rosette. Rosette, first presented in [22] by Emina Torlak, is a framework for automatically enabling solver-aided queries in domain-specific languages (DSLs). Rosette compiles a subset of Racket into constraints (although the full language can be reduced to this subset), thereby enabling the same for any DSL embedded in Racket. These constraints are then supplied to a solver, which can facilitate fault-localization, verification, angelic non-determinism, and synthesis. To allow DSLs to interact with the solver directly, Rosette includes first-class *symbolic* constants, which represent holes whose values are constrained by assertions and determined by calls to the solver.

Rosette's synthesis procedure fills a hole in a partially-complete program given a bounded-depth grammar (a "sketch") from which solutions are drawn. Correctness specifications come in the form of assertions or another program that implements the desired functionality. Rosette compiles the grammar into a nondeterministic program capturing every possible AST that the grammar might produce. Nondeterministic choices are compiled as symbolic values. The solver determines the

nondeterministic choices necessary to fill the hole in a way that conforms to the specification, thereby selecting an appropriate AST and solving the synthesis problem.

In [23], Torlak details the symbolic compiler underlying Rosette. The paper includes a textbook-quality discussion of two possible techniques for compiling Racket into constraints: *symbolic execution* and *bounded model checking*. In symbolic execution, the compiler must individually explore every possible code path of the program, branching on loops and conditionals. This approach concretely executes as much of the program as possible but produces an exponential-sized output. In contrast, bounded model checking joins program states back together after branching, creating new symbolic ϕ-values representing values merged from different states. The resulting formula is smaller but contains fewer concrete values. Both approaches require bounded-depth loop unrolling and function inlining to ensure termination of the compilation procedure.

Torlak adopts a hybrid strategy that draws on the strengths of both techniques: Rosette uses bounded model checking but, when merging, first joins values like lists structurally rather than symbolically, preserving opportunities for concrete evaluation. The compiler assumes that the caller has implemented some form of termination guarantee, such as structural recursion on an inductive datastructure or an explicit limit on recursion depth.

Kaplan, Leon, and related synthesis work. Viktor Kuncak's Kaplan [10] explores the design space of languages with first-class constraints, symbolic values, and the ability to make calls to a solver at run-time in a manner similar to Rosette, adding these features to a pure subset of Scala. Kaplan's constraints are functions from inputs parameters to a boolean expression. Calling *solve* on a constraint generates concrete values for the input parameters that cause the boolean expression to be true. Kaplan also includes mechanisms for iterating over all solutions to a constraint, choosing a particular solution given optimality criteria, or lazily extracting symbolic rather than concrete results. Since constraints are simply functions, they can be composed or combined arbitrarily.

Kaplan introduces a new control-flow construct: an `assuming`...`otherwise` block. This structure functions like an `if`-statement where the guard is a constraint rather than a boolean expression. If the constraint has a solution, the code in the assuming branch is executed; if it is unsatisfiable, the `otherwise` block executes instead.

Kuncak has also explored the opposite approach, performing synthesis entirely at compile-time rather than run-time. Invoking a solver at run-time, he argues, is unpredictable and unwieldy. Ideally, synthesis should function as a compiler service that provides immediate, static feedback to developers and always succeeds when a synthesis problem is feasible. In [11], Kuncak attempts to turn decision procedures for various theories (i.e., linear arithmetic and sets) into corresponding, *complete* synthesis procedures. As in Kaplan, users specify constraints, but the synthesis engine derives a program that statically produces values for the constraint's inputs instead of invoking a solver at run-time.

Combining the lessons of these two systems, Kuncak developed a synthesis engine for creating general-purpose, recursive programs [8]. The synthesizer uses a hybrid algorithm with deductive and inductive steps. The synthesis problem is expressed as a logical specification to which a program must conform rather than individual input-output examples.

In the deductive step, the synthesizer applies various rules for converting a specification into a program in a fashion similar to type-directed synthesis. Special rules are generated to facilitate recursively traversing each available inductive datatype.

At the leaves of this deductive process, the synthesizer uses an inductive CEGIS step to choose candidate expressions. Like Rosette, it encodes bounded-depth ASTs that generate terms of a par-

ticular type as nondeterministic programs. By querying a solver for the nondeterministic choices that satisfy the specification, the synthesizer derives a candidate expression. The synthesizer then attempts to verify this expression. If verification fails, it receives a new counterexample that it integrates into the expression-selection process. The synthesizer increases the AST depth until some maximum value, at which point it determines that the synthesis problem is infeasible. As the program size grows, verification becomes expensive, so the synthesizer uses concrete execution to eliminate candidate programs more quickly.

One notable innovation behind this system is that of *abductive reasoning*, where conditionals that partition the problem space are synthesized, creating more restrictive path conditions that make it easier to find candidate solutions. This method is similar to that in Escher [2], which creates a conditional whenever it synthesizes a program that partially satisfies the input-output examples.

Sketch. The synthesis technique of sketching (introduced in Section 5.1 with the StreamBit language) eventually culminated in a C-like, general-purpose programming language, called Sketch, which synthesizes values for integer and boolean holes. Although Sketch's holes are limited to these particular types, users can creatively place holes in conditional statements to force the synthesizer to choose between alternative code blocks in a manner similar to Rosette. Like StreamBit, a Sketch specification is a program that implements the desired functionality.

Sketch was first introduced by Solar-Lezama in [20] and continues to see active development [12]. In [20], Solar-Lezama describes both the core Sketch language and the underlying, solver-aided synthesis procedure, CEGIS. Given a specification program s and a partial program p with input values $\vec{x}$ and holes $\vec{h}$, Sketch's synthesis problem may be framed as the question of whether $\exists\ \vec{h},\ \forall\ \vec{x},\ p(\vec{h},\vec{x}) = s(\vec{x})$. Solving this quantified boolean formula is Σ_2-complete (i.e., at least as hard as NP-complete), making it too difficult to handle directly.

Instead, Solar-Lezama presents a two-step process for synthesizing a solution. First, a solver guesses values $\vec{v}$ for $\vec{h}$. A verifier then checks whether $\forall\ \vec{x},\ p(\vec{v},\vec{x}) = s(\vec{x})$. If the verification step succeeds, then $\vec{v}$ is a solution to the synthesis problem; otherwise, the verifier returns a counterexample input $\vec{c_1}$. The solver then guesses new values $\vec{v}$ for $\vec{h}$ by integrating $\vec{c_1}$ into the formula: $\exists\ \vec{h},\ p(\vec{h},\vec{c_1}) = s(\vec{c_1})$. If no values for $\vec{h}$ can be found, synthesis fails; otherwise, we return to the verification step and the CEGIS loop continues. In general, it takes no more than a few examples to pinpoint values for $\vec{h}$.

The Sketch language presented in the original paper includes several macro-like operators for managing holes at a higher level. There are *loop* statements that repeat a hole for a particular number of iterations and *generators* that abstract holes into recursive hole-generating functions. Both structures are evaluated and inlined at compile-time. Later versions of Sketch introduce a regular expression-like syntax for describing AST grammars and assertions in place of a specification program [13]. Additional papers add features like concurrency [21]. Solar-Lezama's dissertation collects these ideas in a single document [13].

Autograder. In an exemplary instance of using synthesis as a service, Rishabh Singh built on Sketch to create an autograder for introductory computer science assignments [18]. The autograder relies on two premises: (1) a solution program is available and (2) common errors are known. A user specifies *rewrite rules* consisting of possible corrections to a student program.

The autograder compiles student programs into Sketch. At each location where a rewrite rule might apply, the autograder inserts a conditional statement guarded with a hole that chooses

whether to leave the previous statement in place or take the result of a rewrite rule. When Sketch is run on the student program with the solution as its specification, the values for the holes determine the corrections necessary to fix the student program. Singh had to modify Sketch's CEGIS algorithm to also consider an optimization constraint: the holes should be filled such that the smallest number of changes are made to the original program.

5.3 Theorem Proving

Frank Pfenning's lecture notes for his course *Automated Theorem Proving* [15] are a textbook-quality resource on techniques for efficient proof search in constructive logic. Building off of an introduction to natural deduction and its more wieldy analog, the sequent calculus, Pfenning describes methods for performing forward and backward search over first order logic. Among strategies presented are *focusing*, in which one should eagerly apply proof steps whose effects are invertible, and the *subformula property*, which states that, in forward search, the only necessary proof steps are those that produce a subformula of the goal statement. The concepts portrayed in Pfenning's text, together with Gentzen's classic treatment of the sequent calculus [6], represent the theoretical underpinnings of many of the ideas contained in this thesis.

6 Future Work

6.1 Additional Language Features

Records and Polymorphism. Although the original type-directed synthesis system included only the simply-typed lambda calculus with small additions, this representation already captured a useful portion of a pure subset of ML. This thesis adds a number of missing components, including tuples and the unit type. The most prominent absent features are records and polymorphism.

Records should generally resemble tuples, however their full ML specification would require adding subtyping to the synthesis system. The primary dividends of integrating records would be to explore synthesis in the presence of subtyping and to help transform the type-directed synthesis implementation from a research prototype into a fully-featured programming tool.

Polymorphism presents a wider array of challenges and benefits. System F-like parametric polymorphism would require significant theoretical, syntactic, and implementation changes to the synthesis system. Adding the feature would, however, bring the synthesis language far closer to ML and dramatically improve performance in some cases.

Consider, for example, the process of synthesizing *map*. If the synthesizer knows that *map* will always be applied to a list of natural numbers, it is free use the `S` constructor or to match on an element of the list. If, instead, *map* were to work on a polymorphic list about whose type we know nothing, then the synthesizer will have far fewer choices and will find the desired program faster. This behavior is a virtuous manifestation of Reynolds' abstraction theorem [17].

Solver-Aided Types. The existing synthesis system works only on algebraic datatypes. A logical next step would be to integrate true, 32-bit integers and 64-bit floating point numbers. Doing so within the existing framework, however, would be impossibly inefficient. We could represent integers as an algebraic datatype with 2^{32} base cases, but the synthesizer would time out while guessing every possible number.

Instead, we might integrate with a solver equipped with the theory of bitvectors. These tools have proven to be quite efficient in other, first-order synthesis systems [13]. By combining the higher-order power of type-directed synthesis with the performance of modern SMT-solvers, we could leverage the best of both approaches while bringing the system's language closer to full ML.

Monads and Linear Types. Type-directed synthesis is far more extensible than most other synthesis frameworks: in order to add a new language feature, we need only consult the corresponding system of logic. A number of desirable features map to well-understood logical frameworks. For example, monads are captured by lax logic and resources by linear logic. By integrating these systems of logic into our synthesis judgments, we could generate effectful programs that raise exceptions or read from and write to files. These operations are at the core of many "real" programs but remain beyond the capabilities of any existing synthesis system.

6.2 Refinements and Intersections

In a general sense, types represent (possibly-infinite) sets of expressions. Although we are used to the everyday notions of integers, floating-point numbers, and boolean values, we could also craft types that contain a single item or even no elements at all. This is the insight behind *refinement types* [5], which create new, more restrictive subtypes out of existing type definitions. For example, we might create the type `Nil` as a subtype of `list` that contains only the empty list and `Cons(nat, list)` that captures all non-empty lists. Refinements may be understood as simple predicates on types: stating that x has refinement `2` is equivalent to the type $\{x \in \texttt{nat} \mid x = 2\}$.

The entire example grammar ε from the type-directed synthesis system comprises a language of refinements on the type system τ. In order to fully express partial functions as refinements, we need one additional operator: intersection. The refinement $\wedge(\varepsilon_1, \varepsilon_2)$ requires that an expression satisfy both ε_1 and ε_2. For example, the `inc` function should obey the both $0 \to 1$ and $1 \to 2$. Partial functions are simply intersections of arrow refinements.

There are two benefits to characterizing examples as types. The first is that, by lifting examples into the type system, we are able to fully exploit the Curry-Howard isomorphism and decades of research into automated theorem proving. Rather than forging ahead into unknown territory with examples, we are treading familiar ground with an especially nuanced type system. The second is that we can freely draw on more complex operators from refinement type research to enrich our example language. For instance, we might add base types as refinements so users can express the set of all non-empty lists as discussed before, or integrate union types to complement intersections.

Still richer refinements are possible. One intriguing avenue is that of negation, a refinement that functions as the set complement operator over types. Negation allows us to capture the concept of negative examples, for example $\mathsf{not}(0 \to 1)$, or to express to the synthesizer that a particular property of a program does not hold. Adding negation to the type-directed synthesis system enables new algorithms for generating programs, including a type-directed variant of CEGIS. Negation represents one of many possible specification language enhancements that refinements make possible. We might consider adding quantifiers over specifications or even dependent refinements.

7 Conclusions

Type-directed synthesis offers a new approach to an old, popular, and challenging problem. Where most existing systems rely on solvers to generate programs from specifications, type-directed synthesis applies techniques from automated theorem proving. At the time of writing, however, type-directed synthesis extends only to a minimal language: the simply-typed lambda calculus with recursive functions and algebraic datatypes.

In this thesis, I have described extending type-directed synthesis to include product types, which required rethinking both the underlying theory and the prototype implementation. Products simplify many of the existing rules but invite a host of other theoretical challenges. I leveraged the fact that they are invertible to speed up the proof search process with focusing, which in turn required new judgment forms and additional proofs about the behavior of the updated rules.

The prototype implementation now efficiently synthesizes tuples and projections, validating the new theory. More importantly, it concretely demonstrates a key advantage of type-directed synthesis: extending the algorithm with new syntax is as simple as integrating the corresponding systems of logic. Although products represent a seemingly small addition, they pave the way for the synthesis of other, more powerful language features.

8 Acknowledgements

I would like to thank:

My adviser, Professor David Walker, who has spent the past two years teaching me nearly everything I know about programming languages and research. I am grateful beyond measure for his patient guidance, generosity with his time, and willingness to share his expertise.

Professor Steve Zdancewic and Peter-Michael Osera at the University of Pennsylvania, who kindly shared their research project and allowed me to run rampant in their codebase. Their advice, insight, and patience made this thesis possible.

My parents and sister, who endured the long days and late nights that it took to bring this research to fruition and kindly read countless drafts of this seemingly incomprehensible paper.

References

[1] Syntax-guided synthesis. URL `http://www.sygus.org/`.

[2] A. Albarghouthi, S. Gulwani, and Z. Kincaid. Recursive program synthesis. In *Computer Aided Verification*, pages 934–950, 2013.

[3] R. Alur, R. Bodik, G. Juniwal, M. M. Martin, M. Raghothaman, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, and A. Udupa. Syntax-guided synthesis. In *Formal Methods in Computer-Aided Design (FMCAD), 2013*, pages 1–17. IEEE, 2013.

[4] R. Bodik, S. Chandra, J. Galenson, D. Kimelman, N. Tung, S. Barman, and C. Rodarmor. Programming with angelic nondeterminism. In *ACM Sigplan Notices*, volume 45, pages 339–352. ACM, 2010.

[5] T. Freeman and F. Pfenning. *Refinement types for ML*, volume 26. ACM, 1991.

[6] G. Gentzen. Investigations into logical deduction. *American philosophical quarterly*, pages 288–306, 1964.

[7] S. Gulwani. Automating string processing in spreadsheets using input-output examples. In *ACM SIGPLAN Notices*, volume 46, pages 317–330. ACM, 2011.

[8] E. Kneuss, I. Kuraj, V. Kuncak, and P. Suter. Synthesis modulo recursive functions. In *Proceedings of the 2013 ACM SIGPLAN international conference on Object oriented programming systems languages & applications*, pages 407–426. ACM, 2013.

[9] E. Kneuss, I. Kuraj, V. Kuncak, and P. Suter. Synthesis modulo recursive functions. In *Proceedings of the 2013 ACM SIGPLAN international conference on Object oriented programming systems languages & applications*, pages 407–426. ACM, 2013.

[10] A. S. Köksal, V. Kuncak, and P. Suter. Constraints as control. In *ACM SIGPLAN Notices*, volume 47, pages 151–164. ACM, 2012.

[11] V. Kuncak, M. Mayer, R. Piskac, and P. Suter. Complete functional synthesis. *ACM Sigplan Notices*, 45(6):316–329, 2010.

[12] A. S. Lezama. Armando solar-lezama. URL `http://people.csail.mit.edu/asolar/`.

[13] A. S. Lezama. *Program synthesis by sketching*. PhD thesis, University of California, Berkeley, 2008.

[14] P.-M. Osera and S. Zdancewic. Type-and-example-directed program synthesis. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM, 2015.

[15] F. Pfenning. Automated theorem proving, 2004. URL `http://www.cs.cmu.edu/~fp/courses/atp/index.html`.

[16] F. Pfenning and R. Davies. A judgmental reconstruction of modal logic. *Mathematical structures in computer science*, 11(04):511–540, 2001.

[17] J. C. Reynolds. Types, abstraction and parametric polymorphism. 1983.

[18] R. Singh, S. Gulwani, and A. Solar-Lezama. Automated feedback generation for introductory programming assignments. In *ACM SIGPLAN Notices*, volume 48, pages 15–26. ACM, 2013.

[19] A. Solar-Lezama, R. Rabbah, R. Bodík, and K. Ebcioğlu. Programming by sketching for bit-streaming programs. *ACM SIGPLAN Notices*, 40(6):281–294, 2005.

[20] A. Solar-Lezama, L. Tancau, R. Bodik, S. Seshia, and V. Saraswat. Combinatorial sketching for finite programs. In *ACM Sigplan Notices*, volume 41, pages 404–415. ACM, 2006.

[21] A. Solar-Lezama, C. G. Jones, and R. Bodik. Sketching concurrent data structures. In *ACM SIGPLAN Notices*, volume 43, pages 136–148. ACM, 2008.

[22] E. Torlak and R. Bodik. Growing solver-aided languages with rosette. In *Proceedings of the 2013 ACM international symposium on new ideas, new paradigms, and reflections on programming & software*, pages 135–152. ACM, 2013.

[23] E. Torlak and R. Bodik. A lightweight symbolic virtual machine for solver-aided host languages. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, page 54. ACM, 2014.

A Proofs

A.1 Notes

In many of the judgments below, the constructor context (Σ) has been elided for the sake of readability. Since none of the synthesis or focusing judgments modify the constructor context, it is assumed to be a fixed and globally available entity.

A.2 Behavior of Focusing

Lemma 2.1: Properties of Focusing

IF $\quad \Sigma_1 \mid \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Sigma_2 \mid \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2 \quad$ THEN $\quad \Sigma_1 = \Sigma_2,\ \Gamma_1 \subseteq \Gamma_2,\ \mathcal{A}_1 \subseteq \mathcal{A}_2$

Proof: Immediately follows from the focusing judgments. $\square$

Theorem 2.2: Progress of Focusing

$\mathcal{F}_1 = \cdot \quad$ OR $\quad \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2$

Proof: By case analysis on $\mathcal{F}_1$.

Case: $\mathcal{F}_1 = \cdot$. $\square$

Case: $\mathcal{F}_1 = \mathcal{F}, E : \tau$. *Proof by case analysis on* τ.

Subcase: $\tau = \mathsf{unit}$.	*Apply* FOCUS-UNIT.	$\square$
Subcase: $\tau = \mathcal{B}$.	*Apply* FOCUS-BASE.	$\square$
Subcase: $\tau = \tau_1 \to \tau_2$.	*Apply* FOCUS-FUN.	$\square$
Subcase: $\tau = \tau_1 * \ldots * \tau_n$.	*Apply* FOCUS-TUPLE.	$\square$

Theorem 2.3: Preservation of Focusing

IF $\quad \vdash \Gamma_1\ wf, \quad \vdash \mathcal{A}_1\ wf, \quad \vdash \mathcal{F}_1\ wf \quad$ AND $\quad \Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2$
THEN $\vdash \Gamma_2\ wf, \quad \vdash \mathcal{A}_2\ wf, \quad \vdash \mathcal{F}_2\ wf$

Proof: By case analysis on the judgment $\Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2$.

Case: FOCUS-UNIT.

(1)	$\mathcal{F}_1 = \mathcal{F}_2, E : \mathsf{unit}$	By FOCUS-UNIT
(2)	$\Gamma_2 = \Gamma_1, E : \mathsf{unit}$	By FOCUS-UNIT
(3)	$\mathcal{A}_1 = \mathcal{A}_2$	By FOCUS-UNIT
(4)	$\vdash \mathcal{F}_1\ wf$	Assumption
(5)	$\vdash \mathcal{A}_1\ wf$	Assumption
(6)	$\vdash \Gamma_1\ wf$	Assumption
(7)	$\vdash \mathcal{F}_2\ wf$	4 and inversion on TYPE-CTX-ONE-WF
(8)	$\vdash E : \mathsf{unit}$	4 and inversion on TYPE-CTX-ONE-WF
(9)	$\vdash \mathcal{A}_2\ wf$	(3) and (5)
(10)	$\vdash \Gamma_2\ wf$	(6), (9), TYPE-CTX-ONE-WF $\square$

Case: FOCUS-BASE, FOCUS-FUN. *Similar.* $\square$

Case: FOCUS-TUPLE.

(1)	$\mathcal{F}_1 = \mathcal{F}_2, E : \tau_1 * ... * \tau_m$	By FOCUS-TUPLE
(2)	$\Gamma_2 = \Gamma_1$	By FOCUS-TUPLE
(3)	$\mathcal{A}_1 = \mathcal{A}_2, \pi_1\ E : \tau_1, ..., \pi_n\ E : \tau_n$	By FOCUS-TUPLE
(4)	$\vdash \mathcal{F}_1\ \ wf$	Assumption
(5)	$\vdash \mathcal{A}_1\ \ wf$	Assumption
(6)	$\vdash \Gamma_1\ \ wf$	Assumption
(7)	$\vdash \mathcal{F}_2\ \ wf$	4 and inversion on TYPE-CTX-ONE-WF
(8)	$\vdash E : \tau_1 * ... * \tau_m$	4 and inversion on TYPE-CTX-ONE-WF
(9)	$\forall i \in n,\ \vdash \pi_i\ E : \tau_i$	(8) and T-PROJ
(10)	$\vdash \mathcal{A}_2\ \ wf$	(5), (9), and TYPE-CTX-ONE-WF
(11)	$\vdash \Gamma_2\ \ wf$	(2) and (6) $\square$

Lemma 2.4: Potential of Focusing

Let $\phi(\tau)$ be the size of the tree comprising the type τ.

Let $\Phi(\mathcal{F}) = \sum_{E_i:\tau_i \in \mathcal{F}} \phi(\tau_i)$.

IF $\Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2$ THEN $\Phi(\mathcal{F}_2) < \Phi(\mathcal{F}_1)$

Proof: By case analysis on the judgment $\Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow \Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \mathcal{W}_2$.

Case: FOCUS-UNIT. *$\mathcal{F}_2 \subset \mathcal{F}_1$, so $\Phi(\mathcal{F}_2) < \Phi(\mathcal{F}_1)$.* $\square$

Case: FOCUS-BASE, FOCUS-FUN. *Similar.* $\square$

Case: FOCUS-TUPLE. *Removes $E : \tau_1 * ... * \tau_m$ from $\mathcal{F}_1$ and adds $\pi_1\ E : \tau_1$, ..., $\pi_n\ E : \tau_n$ to $\mathcal{F}_1$. $\sum_{i \in n} \phi(\tau_i) < \phi(\tau_1 * ... * \tau_m)$, implying that $\Phi(\mathcal{F}_2) < \Phi(\mathcal{F}_1)$.* $\square$

Theorem 2.5: Termination of Focusing

$\Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow^* \Gamma_2; \mathcal{A}_2; \cdot \mid \mathcal{W}_2$

Proof: By the potential method and case analyis on $\mathcal{F}_1$.

Case: $\mathcal{F}_1 = \cdot$. Apply FOCUS-CLOSURE-BASE. $\square$

Case: $\mathcal{F}_1 = \mathcal{F}, E : \tau$. By Theorem 2.2, we can apply a focusing judgment to any non-empty $\mathcal{F}$. Therefore, apply FOCUS-CLOSURE-STEP. *By Lemma 2.4, $\Phi(\mathcal{F})$ decreases monotonically at each such step. $\Phi(\cdot) = 0$, so this process eventually terminates with $\mathcal{F}_1 = \cdot$.*

A.3 Order of Focusing in the Proof Search Process

Theorem 3.1: Focusing is Deterministic

IF $\Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow^* \Gamma_{2a}; \mathcal{A}_{2a}; \cdot \mid \mathcal{W}_2$ AND $\Gamma_1; \mathcal{A}_1; \mathcal{F}_1 \mid \mathcal{W}_1 \Longrightarrow^* \Gamma_{2b}; \mathcal{A}_{2b}; \cdot \mid \mathcal{W}_2$
THEN $\Gamma_{2a} = \Gamma_{2b}$ AND $\mathcal{A}_{2a} = \mathcal{A}_{2b}$

Proof: Immediately follows from the focusing judgments. $\square$

Theorem 3.2: Focusing Need Not Occur Before Synthesis

DEFINE THE JUDGMENTS $\Gamma; \mathcal{A}; \mathcal{F} \vdash \tau \overset{E'}{\rightsquigarrow} E$ AND $\Gamma; \mathcal{A}; \mathcal{F} \mid \mathcal{W} \vdash \tau \overset{I'}{\rightsquigarrow} I$ TO BE IDENTICAL TO THOSE IN FIGURE 7 EXCEPT THAT $\mathcal{F}$ NEED NOT BE EMPTY FOR ANY SYNTHESIS RULE TO BE APPLIED.

$\Gamma; \mathcal{A}; \mathcal{F} \mid \mathcal{W} \vdash \tau \overset{I}{\rightsquigarrow} I$ IFF $\Gamma; \mathcal{A}; \mathcal{F} \mid \mathcal{W} \vdash \tau \overset{I'}{\rightsquigarrow} I$

$\Gamma; \mathcal{A}; \mathcal{F} \vdash \tau \overset{E}{\rightsquigarrow} E$ IFF $\Gamma; \mathcal{A}; \mathcal{F} \vdash \tau \overset{E'}{\rightsquigarrow} E$

Proof: Rather than focusing eagerly, we merely focus lazily as focused terms become necessary. Since focusing is deterministic (Theorem 3.1), the time at which we perform focusing does not affect the expressions we can synthesize. □

A.4 Admissibility of Focusing

DEFINE THE JUDGMENTS $\Gamma \vdash \tau \overset{E}{\rightsquigarrow} E$ AND $\Gamma \mid \mathcal{W} \vdash \tau \overset{I}{\rightsquigarrow} I$ TO BE IDENTICAL TO THOSE IN FIGURE 7 WITHOUT FOCUSING. INSTEAD OF FOCUSING TO PROJECT ON TUPLES, WE INTRODUCE THE EGUESS-PROJ RULE AS BELOW:

$$\frac{\Gamma, E_1 : \tau_1 * ... * \tau_m,\ \pi_1\ E_1 : \tau_1,\ ...,\ \pi_m\ E_1 : \tau_m \vdash \tau \overset{E}{\rightsquigarrow} E}{\Gamma, E_1 : \tau_1 * ... * \tau_m \vdash \tau \overset{E}{\rightsquigarrow} E}\ \textsc{EGuess-Proj}$$

Theorem 4.1: Soundness of Focusing

IF $\Gamma; \cdot; \cdot \mid \cdot \vdash \tau \overset{I}{\rightsquigarrow} I$ THEN $\Gamma \mid \cdot \vdash \tau \overset{I}{\rightsquigarrow} I$

IF $\Gamma; \cdot; \cdot \vdash \tau \overset{E}{\rightsquigarrow} I$ THEN $\Gamma \vdash \tau \overset{E}{\rightsquigarrow} I$

Proof: We can transform any derivation of a judgment in the system with focusing into the system without focusing.

Where elimination forms of unit, base, or function type are placed in $\mathcal{F}$ and subsequently moved into Γ using IREFINE-FOCUS *and* EGUESS-FOCUS*, simply remove the focusing judgments and move the elimination forms directly into Γ.*

Where elimination forms of tuple type are placed in $\mathcal{F}$ and projected upon into their non-tuple constituents using IREFINE-FOCUS *and* EGUESS-FOCUS*, insert uses of the* EGUESS-PROJ *judgment that perform the same projections.* □

Theorem 4.2: Completeness of Focusing

IF $\Gamma \mid \cdot \vdash \tau \overset{I}{\rightsquigarrow} I$ THEN $\Gamma; \cdot; \cdot \mid \cdot \vdash \tau \overset{I}{\rightsquigarrow} I$

IF $\Gamma \vdash \tau \overset{E}{\rightsquigarrow} I$ THEN $\Gamma; \cdot; \cdot \vdash \tau \overset{E}{\rightsquigarrow} I$

Proof: We can transform any derivation of a judgment in the system without focusing into the system with focusing.

Whenever an elimination form is inserted into Γ, instead insert it into $\mathcal{F}$ and perform focusing.

Elimination forms of unit, base, or function type may be used exactly as before aside from additional bookkeeping to move them from $\mathcal{F}$ to Γ.

Even in the system without focusing, tuples must still be fully projected before they may be used as a result of the restriction that we can only synthesize programs that are eta-long. EGUESS-CTX *only moves terms from the context to goal position at base type. No elimination form judgment other than* EGUESS-PROJ *may make use of tuples. Therefore, any tuple placed into the context in the system without focusing is either unused or fully projected. Focusing merely performs this projection eagerly.* □

A.5 Soundness of Synthesis Judgments

Theorem 5.1: Soundness of Synthesis Judgments

IF $\Gamma; \mathcal{A}; \mathcal{F} \mid \overline{\langle \mathcal{E}; \varepsilon \rangle}^{n} \vdash \tau \overset{I}{\rightsquigarrow} I$ THEN $[\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash I : \tau$ AND $\forall i \in n,\ \mathcal{E}_i(I) \rightarrow^* \varepsilon_i$

IF $\Gamma; \mathcal{A}; \mathcal{F} \vdash \tau \overset{E}{\rightsquigarrow} E$ THEN $[\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash E : \tau$

Proof: By simultaneous induction on the IREFINE *and* EGUESS *judgments.*

Case EGUESS-CTX.

(1) $E : \tau \in \Gamma$ Inversion $\square$

Case EGUESS-APP.

(1) $\Gamma, E_1 : \tau_1 \to \tau_2; \mathcal{A}; E_1\ I_1 : \tau_2 \vdash \tau \overset{E}{\rightsquigarrow} E$ Inversion
(2) $[\Gamma, E_1 : \tau_1 \to \tau_2, \mathcal{A}, E_1\ I_1 : \tau_2]_{vars} \vdash e : \tau$ Inductive Hypothesis
(3) $[\Gamma, E_1 : \tau_1 \to \tau_2, \mathcal{A}]_{vars} \vdash e : \tau$ Vars relation $\square$

Case EGUESS-FOCUS.

(1) $[\Gamma_1, \mathcal{A}_1, \mathcal{F}_1]_{vars} = [\Gamma_2, \mathcal{A}_2, \mathcal{F}_2]_{vars}$ Immediate from focusing judgments
(2) $\Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \vdash \tau \overset{E}{\rightsquigarrow} E$ Inversion
(3) $[\Gamma_2, \mathcal{A}_2, \mathcal{F}_2]_{vars} \vdash E : \tau$ Inductive Hypothesis
(4) $[\Gamma_1, \mathcal{A}_1, \mathcal{F}_1]_{vars} \vdash E : \tau$ (1) and (3) $\square$

Case IREFINE-FOCUS.

(1) $[\Gamma_1, \mathcal{A}_1, \mathcal{F}_1]_{vars} = [\Gamma_2, \mathcal{A}_2, \mathcal{F}_2]_{vars}$ Immediate from focusing judgments
(2) $\Gamma_2; \mathcal{A}_2; \mathcal{F}_2 \mid \overline{\langle \mathcal{E}; \varepsilon \rangle}^{n} \vdash \tau \overset{I}{\rightsquigarrow} E$ Inversion
(3) $[\Gamma_2, \mathcal{A}_2, \mathcal{F}_2]_{vars} \vdash E : \tau$ Inductive Hypothesis
(4) $[\Gamma_1, \mathcal{A}_1, \mathcal{F}_1]_{vars} \vdash E : \tau$ (1) and (3)
(5) $\forall i \in n,\ \mathcal{E}_i(I) \to^* \varepsilon_i$ Inversion $\square$

Case IREFINE-GUESS.

(1) $\Gamma; \mathcal{A}; \cdot \vdash \mathcal{B} \overset{E}{\rightsquigarrow} E$ Inversion
(2) $[\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash E : \mathcal{B}$ Inductive Hypothesis
(3) $\forall i \in n,\ \mathcal{E}_i(E) \to^* \varepsilon_i$ Inversion $\square$

Case IREFINE-UNIT.

(1) $[\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash () : \mathsf{unit}$ T-UNIT
(2) $\forall i \in n,\ \mathcal{E}_i(()) \to^* ()$ S-CLOSURE-BASE $\square$

Case IREFINE-CTOR.

(1) $C : \tau \to \mathcal{B} \in \Sigma$ Inversion
(2) $\Gamma; \mathcal{A}; \cdot \mid \overline{\langle \mathcal{E}_i; I_i \rangle}^{i \in n} \vdash \tau \overset{I}{\rightsquigarrow} I$ Inversion
(3) $[\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash I : \tau$ Inductive Hypothesis
(4) $[\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash C\ I : \mathcal{B}$ T-CTOR, (1), (3)
(5) $\forall i \in n,\ \mathcal{E}_i(I) \to^* I_i$ Inductive Hypothesis
(6) $\forall i \in n,\ \mathcal{E}_i(C\ I) \to^* C\ I_i$ S-CTOR, (5) $\square$

Case IREFINE-TUPLE.

(1) $\forall i \in m,\ \Gamma; \mathcal{A}; \cdot \mid \overline{\langle \mathcal{E}_j; I_{(i,j)} \rangle}^{j \in n} \vdash \tau_i \overset{I}{\rightsquigarrow} I_i$ Inversion
(2) $\forall i \in m, [\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash I_i : \tau_i$ Inductive Hypothesis
(3) $[\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash (I_1, ..., I_m) : \tau_1 * ... * \tau_m$ T-TUPLE, (2)
(4) $\forall (i,j) \in (m,n),\ \mathcal{E}_j(I_i) \to^* I_{(i,j)}$ Inductive Hypothesis
(5) $\forall j \in n,\ \mathcal{E}_j((I_1, ..., I_m)) \to^* (I_{(1,j)}, ..., I_{(m,j)})$ S-TUPLE, (4) $\square$

Case IREFINE-MATCH.

(1) $\Gamma; \mathcal{A}; \cdot \vdash \mathcal{B} \overset{E}{\rightsquigarrow} E$ Inversion
(2) $\forall j \in m,\ C_j : \tau_j \to \mathcal{B} \in \Sigma$ Inversion
(3) $\forall j \in m,\ \Sigma \mid \Gamma; \mathcal{A}; x : \tau_j \mid \{ \langle \mathcal{E}_i, x \triangleright v; \varepsilon_i \rangle \mid \mathcal{E}_i(E) \to^* C_j\ v \} \vdash \tau \overset{I}{\rightsquigarrow} I_j$ Inversion
(4) $[\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash E : \mathcal{B}$ Inductive Hypothesis, (1)
(5) $\forall j \in m,\ [\Gamma, \mathcal{A}, \mathcal{F}, x : \tau_j]_{vars} \vdash I_j : \tau$ Inductive Hypothesis, (2)
(6) $[\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash \mathsf{match}\ E\ \mathsf{with}\ \overline{C_j\ x \to I_j}^{j \in m} : \tau$ T-MATCH, (2), (4), (5)
(7) $\forall i \in n,\ \mathcal{E}_i(E) \to^* C_j\ v$ Inversion
(8) $\forall i \in n,\ \text{IF } \mathcal{E}_i(E) \to^* C_j\ v \text{ THEN } \mathcal{E}_i, x \triangleright v(I_j) \to^* \varepsilon_i$ Inductive Hypothesis, (3)
(9) $\forall i \in n,\ \mathcal{E}_i(\mathsf{match}\ E\ \mathsf{with}\ \overline{C_j\ x \to I_j}^{j \in m}) \to^* \varepsilon_i$ S-MATCH1, S-MATCH2, (7), (8) $\square$

Case IREFINE-FIX.

(1) $\Gamma; \mathcal{A}; f : \tau_1 \to \tau_2, x : \tau_1 \mid \overline{\langle \mathcal{E}_i, f \triangleright \overline{v_{(i,j)} \Rightarrow \varepsilon_{(i,j)}}^{j \in m_i}, x \triangleright v_{(i,k)}; \varepsilon_{(i,k)} \rangle}^{(i,k) \in (n,m_i)} \vdash \tau_2 \overset{I}{\rightsquigarrow} I$ Inversion

(2) $[\Gamma, \mathcal{A}, \mathcal{F}, f : \tau_1 \to \tau_2, x : \tau_1]_{vars} \vdash I : \tau_2$ Inductive Hypothesis

(3) $[\Gamma, \mathcal{A}, \mathcal{F}]_{vars} \vdash \mathsf{fix}\, f(x : \tau_1) : \tau_2 = I : \tau_1 \to \tau_2$ T-ABS, (2)

(4) $\forall (i,k) \in (n, m_i),\ \mathcal{E}_i, f \triangleright \overline{v_{(i,j)} \Rightarrow \varepsilon_{(i,j)}}^{j \in m_i}, x \triangleright v_{(i,k)}(I) \to^* \varepsilon_{(i,k)}$ Inductive Hypothesis

(5) $\forall i \in n,\ \mathcal{E}_i(\mathsf{fix}\, f(x : \tau_1) : \tau_2 = I) \to^* \overline{v_{(i,j)} \Rightarrow \varepsilon_{(i,j)}}^{j \in m_i}$ S-CLOSURE-BASE, (4) □

Note that we assume that the fixpoint and the partial function are equal to one another since they have the same behavior on all inputs for which the partial function is defined.